



Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

UTILITY PATENT APPLICATION TRANSMITTAL

(Only for new nonprovisional applications under 37 CFR § 1.53(b))

Attorney Docket No.	1848.0060001
First Inventor or Applicant Identifier	Mark Scott
Title	Method, System, and Computer Program Product for Managing Database Servers and Services
Express Mail Label No.	

APPLICATION ELEMENTS

See MPEP chapter 600 concerning utility patent application contents.

ADDRESS TO

 Assistant Commissioner for Patents
 Box Patent Application
 Washington, DC 20231

1. ☒ * Fee Transmittal Form (e.g., PTO/SB/17)
(Submit an original, and a duplicate for fee processing)
2. ☒ Specification [Total Pages 282]
(preferred arrangement set forth below)
 - Descriptive title of the invention
 - Cross Reference to Related Applications
 - Background of the Invention
 - Brief Summary of the Invention
 - Brief Description of the Drawings (if filed)
 - Detail of Description
 - Claim(s)
 - Abstract of the Disclosure
3. ☒ Drawing(s) (35 U.S.C. 113) [Total Sheets 117]
4. ☒ Unexecuted Oath or Declaration [Total Pages 4]
 a. ☐ Newly executed (original or copy)
 b. ☐ Copy from a prior application (37 CFR 1.63(d)) (for continuation/divisional with Box 17 completed)
 [Note Box 5 below]
 i. ☐ **DELETION OF INVENTOR(S)**
 Signed statement attached deleting inventor(s) named in the prior application, see 37 CFR §§ 1.63(d)(2) and 1.33(b).
5. ☐ Incorporation By Reference (available if Box 4b is checked)
 The entire disclosure of the prior application, from which a copy of the oath or declaration is supplied under Box 4b, is considered as being part of the disclosure of the accompanying application and is hereby incorporated by reference therein.

6. ☐ Microfiche Computer Program (Appendix)
7. Nucleotide and/or Amino Acid Sequence Submission (if applicable, all necessary)
 a. ☐ Computer Readable Copy
 b. ☐ Paper Copy (identical to computer copy)
 c. ☐ Statement verifying identity of above copies

ACCOMPANYING APPLICATION PARTS

8. ☐ Assignment Papers (cover sheet & document(s))
9. ☐ 37 CFR 3.73(b) Statement (when there is an assignee) ☐ Power of Attorney
10. ☐ English Translation Document (if applicable)
11. ☐ Information Disclosure Statement (IDS)/PTO-1449 ☐ Copies of IDS Citations
12. ☐ Preliminary Amendment
13. ☒ Return Receipt Postcard (MPEP 503)
(Should be specifically itemized)
14. ☐ *Small Entity Statement(s) ☐ Statement filed in prior application, Status still proper and desired
(PTO/SB-08-12)
15. ☐ Certified Copy of Priority Document(s)
(if foreign priority is claimed)
16. ☒ Other 37 C.F.R. § 1.136(a)(3) Authorization
☐ Other
- *NOTE FOR ITEMS 1 & 14: IN ORDER TO BE ENTITLED TO PAY SMALL ENTITY FEES, A SMALL ENTITY STATEMENT IS REQUIRED (37 C.F.R. § 1.29), EXCEPT IF ONE FILED IN A PRIOR APPLICATION IS RELIED UPON (37 C.F.R. § 1.29)

17. If a CONTINUING APPLICATION, check appropriate box, and supply the requisite information below and in a preliminary amendment.

☐ Continuation ☐ Divisional ☒ Continuation-in-Part (CIP) of prior application No: 09/393,658

 Prior application information: Examiner To be assigned Group/Art Unit: To be assigned

18. CORRESPONDENCE ADDRESS

☐ Customer Number
 or Bar Code Label

(Insert Customer No. or Attach bar code label here)

☒ Correspondence
 address below

NAME	STERNE, KESSLER, GOLDSTEIN & FOX P.L.L.C.				
	Attorneys at Law				
ADDRESS	Suite 600, 1100 New York Avenue, N.W.				
CITY	Washington	STATE	DC	ZIP CODE	20005-3934
COUNTRY	USA	TELEPHONE	(202) 371-2600	FAX	(202) 371-2540

NAME (Print/Type)	Michael W. Messenger	Registration No. (Attorney/Agent)	37,575
SIGNATURE		DATE	3/17/00

Burden Hour Statement: this form is estimated to take 0.2 hours to complete. Time will vary depending upon the needs of the individual case. Any comments on the amount of time you are required to complete this form should be sent to the Chief Information Officer, Patent and Trademark Office, Washington, DC 20231. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Assistant Commissioner for Patents, Washington, DC 20231.

STERNE, KESSLER, GOLDSTEIN & FOX P.L.L.C.

ATTORNEYS AT LAW

1100 NEW YORK AVENUE, N.W., SUITE 600
WASHINGTON, D.C. 20005-3934

(202) 371-2600

FACSIMILE (202) 371-2540, (202) 371-6566

jc644 U.S. PTO



03/17/00

ROBERT GREENE STERNE
EDWARD J. KESSLER
JONSE A. GOLDSTEIN
SAMUEL L. FOX
DAVID K.S. CORNWELL
ROBERT W. ESMOND
TRACY-GENE G. DURKIN
MICHELE A. CIMBALA
MICHAEL B. RAY
ROBERT E. SOKOHL
ERIC K. STEFFE
MICHAEL O. LEE

STEVEN R. LUDWIG
JOHN M. COVERT**
LINDA E. ALCORN
RAZ E. FLESHNER
ROBERT C. MILLONIG
MICHAEL V. MESSINGER
JUDITH U. KIM
TIMOTHY J. SHEA, JR.
DONALD R. MCPHAIL
PATRICK E. GARRETT
STEPHEN G. WHITESIDE
JEFFREY T. HELVEY*

HEIDI L. KRAUS
JEFFREY R. KURIN
RAYMOND MILLIEN
PATRICK D. O'BRIEN
LAWRENCE B. BUGAISKY
CRYSTAL D. SAYLES**
EDWARD W. YEE
ALBERT L. FERRO*
DONALD R. BANOWIT
PETER A. JACKMAN
MOLLY A. MCCALL

TERESA U. MEDLER
JEFFREY S. WEAVER
KRISTIN K. VEDOVICH
KENDRICK P. PATTERSON
DONALD J. FEATHERSTONE
GRANT E. REED
VINCENT L. CAPUANO
JOHN A. HARROUN*
MATTHEW M. CATLETT*
NATHAN K. KELLEY**
ALBERT J. FASULO II *

KAREN R. MARKOWICZ**
SUZANNE C. ZISRA**
BRIAN J. DEL BUONO**
ANDREA J. KAMAGE**
NANCY J. LEITH**
TARJA H. NAUKKARINEN**

*BAR OTHER THAN D.C.
**REGISTERED PATENT AGENTS

March 17, 2000

WRITER'S DIRECT NUMBER:

(202) 371-2667

INTERNET ADDRESS:

MIKEM@SKGF.COM

Assistant Commissioner for Patents
Washington, D.C. 20231

Box Patent Application

Re: U.S. Continuation-In-Part Utility Patent Application under 37 C.F.R. § 1.53(b)
Appl. No. To be assigned; Filed: March 17, 2000

For: **Method, System, and Computer Program Product for Managing
Database Servers and Services**

Inventors: Mark Scott, William Wong, Anita Cheng, Simon Ho,
George Irimescu, Dorel Voineag, Min Yao, and Row J. Zadeh

Our Ref: 1848.0060001

Sir:

The following documents are forwarded herewith for appropriate action by the U.S.
Patent and Trademark Office:

1. U.S. Utility Patent Application entitled:

**Method, System, and Computer Program Product for Managing Database
Servers and Services**

and naming as inventors:

**Mark Scott, William Wong, Anita Cheng, Simon Ho, George Irimescu, Dorel
Voineag, Min Yao, and Row J. Zadeh**

the application comprising:

- a. specification containing:
 - i. 255 pages of description prior to the claims;


Assistant Commissioner for Patents
March 17, 2000
Page 2

- ii. 26 pages of claims (91 claims);
- iii. a one (1) page abstract;
- b. One-hundred and seventeen sheets of drawings: (Figures 1A, 1B, 2, 3A-D, 4, 5, 6A, 6B, 7-20, 21A, 21B, 22A-C, 23-48, 49A, 49B, 50-159, 160A, 160B, 161-164);
- 2. USPTO Utility Patent Application Transmittal Form PTO/SB/05;
- 3. Fee Transmittal (Form PTO/SB/17) (in duplicate);
- 4. 37 C.F.R. § 1.136(a)(3) Authorization to Treat a Reply As Incorporating An Extension of Time (in duplicate);
- 5. An unexecuted copy of the Combined Declaration for Patent Application and Power of Attorney;
- 6. Our Check No. 26958 in the amount of \$2,670.00 to cover:
 - \$ 690.00 Basic filing fee (37 C.F.R. § 1.16(a));
 - \$ 702.00 Fees for independent claims over three (37 C.F.R. § 1.16(b));
 - \$1,278.00 Fees for excess claims over twenty (37 C.F.R. § 1.16(c)); and
- 7. Two (2) return postcards.

It is respectfully requested that, of the two attached postcards, one be stamped with the filing date of these documents and returned to our courier, and the other, prepaid postcard, be stamped with the filing date and unofficial application number and returned as soon as possible.

Respectfully submitted,

STERNE, KESSLER, GOLDSTEIN & FOX P.L.L.C.

Per  36,013
Michael V. Messinger
Attorney for Applicants
Registration No. 37,575

Method, System, and Computer Program Product for Managing Database Servers and Services

Inventors: Mark Scott
William Wong
Anita Cheng
Simon Ho
George Irimescu
Dorel Voineag
Min Yao
Row J. Zadeh

Background of the Invention

Cross-Reference to Related Applications

This application is a continuation-in-part of the commonly owned, co-pending U.S. Appl. No. 09/393,658, filed September 10, 1999 (incorporated by reference herein in its entirety).

This application is related to commonly owned, co-pending U.S. patent application entitled "Method, System and Computer Program Product for Managing Routing Servers and Services" filed concurrently herewith (incorporated by reference herein in its entirety).

Field of the Invention

The present invention pertains to voice over the Internet technology.

Related Art

a. Voice over the Internet

Worldwide deregulation and packet-switched technology have brought dramatic changes to the telecommunications industry. Voice communication over packet-switched networks, such as, the Internet, is now in demand. Voice over the Internet technology is now available which integrates a public switched telephone network (PSTN) and the Internet. Voice over the Internet technology is also referred to as, Voice over Internet Protocol (VoIP), Voice over Network (VON), Internet Telephony, and Telephony over the Internet. See, e.g., G. Held, *Voice over Data Networks*, Mc-Graw-Hill Series on Computer Communications (McGraw-Hill Publ.: 1998), entire book; and T. Sheldon, *Encyclopedia of Networking*, Electronic Edition, (Mcgraw-Hill Publ.: 1998), pp.1047-1050. Both of these references are incorporated in their entirety herein by reference. A PSTN is also referred to as a plain old telephone system (POTS). A PSTN can include transmission components (links), switching components (nodes), and billing facilities. See, e.g., Held, chapter 4, "Understanding Voice" and chapter 5, "Telephone Operations," and Sheldon at pp. 959-968. The Internet can be any network or combination of networks that support packet-switching. Such a network can include, but is not limited to, a network supporting the Internet Protocol (IP) and related protocols. See, e.g., Held, chapter 2, "IP and Related Protocols," and Sheldon, pp.503-543.

VoIP systems allow a call originating from a PSTN to be carried over the Internet. In this way, voice traffic can be carried over the Internet bypassing more expensive long-distance carriers. For example, Toll Bypass Carrier services, multi-site corporations, and other users use VoIP systems to reduce the cost of long distance telephone calls to a fraction of one cent per minute.

Vendors provide a product or family of products to support VoIP. For example, Array Telecom Corp. of Herndon, VA, a wholly-owned subsidiary of Comdial Company, has provided an early generation of products that provide a multi-location, turn-key solution to VoIP (Array Tel Series 1000 and 2000, also

called versions 1 and 2). See also B. McConnell, *How to Build an Internet PBX*, <http://www.hellodirect.net/tutorial/ip-phone.htm>, downloaded August 16, 1999.

b. How Voice over IP Telephony Works

5 A VoIP system typically relies on gateways to provide an interface between PSTN and IP networks. FIGs. 1A and 1B illustrate how a call is placed over the Internet. As shown in FIG. 1A, a user calls picks up a telephone 100 and dials a phone number of a phone 140. The called party at phone 140 answers. A connection is then established through gateways 110 and 130 over IP network 120. Voice, fax, and data transmission is then carried out between phones 110 and 140 through this connection. In this way, an expensive long-distance carrier on a PSTN can be bypassed.

10
15
20
25
FIG. 1B shows the operation of a conventional VoIP system in greater detail. A local PBX (private branch exchange) system 105 places a call originating from phone 100 over a PSTN to a local originating gateway 110. Note gateways 110 and 130 are referred to herein as originating and destination gateways for convenience only. In practice, full-duplex communication is carried out and either phone can originate or receive a call. Also, PBX system 105 is programmed to automatically (and selectively) route certain calls from phone 100 to gateway 110. In this way, a user can dial the destination phone number directly in one step. Alternatively, if no PBX or other special preconfiguration at a local exchange carrier (LEC) is performed, a two step process is used where a user calls gateway 110 directly and then in response to a prompt from gateway 110 enters the destination phone number of phone 140.

Next, gateway 110 contacts destination gateway 130 over IP network 120. Destination gateway 130 dials the destination phone number over a PSTN line to the destination phone 140. Gateway 110 then packetizes the voice stream from the PSTN and routes it over IP network 120 to gateway 130. Destination gateway

130 converts the data packets into a voice stream and sends it over a PSTN line to the called party at phone 140. By strategically locating gateways, such as gateways 110 and 130, in desired calling areas long-distance costs are substantially reduced. Indeed, a long-distance call is essentially replaced by two local phone calls and a low cost IP connection.

c. Shortcomings

Even though VoIP has created an enormous enthusiasm, the implementation and success of VoIP networks has been hindered due to at least following four (4) major shortcomings:

- ▶ voice jitter caused by unpredictable bottlenecks of the Internet;
- ▶ unacceptable voice quality due to poor compression techniques;
- ▶ absence of advanced applications to offer services such as prepaid calling or carrier class global call routing; and
- ▶ absence of switching functions such as trunk-to-trunk connections and PBX interfaces.

Summary of the Invention

A method, system, and computer program product that provides gateway, routing and database servers with services for voice over the Internet communication (VoIP) is disclosed herein. In one embodiment, the present invention is a VoIP system that includes a gateway server, a routing server, and a database server. The gateway server handles calls received from a public switched telephone network and a packet-switched network. Messages can be sent between the gateway server, routing server, and database server over the packet-switched network. In another embodiment, the present invention is a

VoIP system that includes an originating gateway server, a terminating gateway server, and a routing server. Both gateway servers handle calls received from a public switched telephone network and a packet-switched network. Messages can be sent between the gateway servers. The routing server is accessed by the gateway servers and provides routing information to the gateway servers over the packet-switched network. In another embodiment, the present invention is a VoIP system that includes an originating gateway server, a terminating gateway server, and a database server. Both gateway servers handle calls received from a public switched telephone network and a packet-switched network. Messages can be sent between the gateway servers. The database server monitors access and records activity on the gateway servers over the packet-switched network.

In addition, a provisioning system is coupled to the database server. A management system is also coupled to the gateway server, routing server, and database server over the packet-switched network.

According to one feature, a network manager automatically queries a client database to determine an update, and sends a message representative of the update to at least one of the gateway server, routing server, database server, and management system over the packet-switched network.

According to another feature of the present invention, a licensing server is used to support licensing of the VoIP system.

In a further embodiment of the present invention, a routing manager component is coupled to each gateway server, and a parser subsystem is coupled to each routing manager.

According to one feature of the present invention, the gateway server and routing server exchange routes using routing models over a packet-switched network.

In one embodiment, the present invention includes a gateway server with a routing manager for connecting bridging calls on a packet-switched network.

The present invention further includes methods for enabling the exchange of route information between gateway server and routing server.

In the present invention, routes can be exported from gateway server and routing server to other gateway server and routing server. Additionally, routes
5 can be transformed by routing servers for use on other gateway servers.

In one embodiment, the present invention includes a route model for local routing and uses local routing to connect gateway servers. The present invention further includes route models for static global routing, disseminated routing, and/or dynamic routing. The present invention can use any number of route
10 models, alone or in combination.

According to another feature, the routes can be either network or telephony based. In addition, the present invention provides for grouping network and telephony routes.

In a further embodiment of the present invention, a routing server is connected over a packet-switched network to a number of gateway servers, and a parser subsystem is included in each routing server. The routing server of the
15 present invention further translates routes based on their origin, use, and/or type.

In one embodiment, the present invention includes routines for determining and prioritizing a call address. Call addresses of the present invention can be stored in routes along with other values (such as access based on the time of day, user account or location) which the gateway servers, routing
20 servers and other systems of the present invention can use.

According to another feature, a database server includes a database access manager, a multi-database client for communicating with any number of database client(s), database client(s) for communicating with database server(s), database
25 server(s) with links to database(s), and database(s) for storing information. The multi-database client provides a middle layer for the relaying of information from any number of connected database clients.

In one embodiment, the present invention includes routines for viewing, configuring and reviewing user accounts, call detail records and event logs. The present invention further includes a management system for configuring component systems.

5 In one embodiment, the present invention includes a graphical user interface that enables a user to configure all the components of the present invention. The present invention further includes a graphical user interface (such as in a screen and panel) for configuring the identification, access, line groups, routing, security, database, management, services, licensing, and hardware setup properties of the components of the system.

10 In a further embodiment, the present invention includes a graphical user interface for controlling the execution of processes on the gateway servers, routing servers and database servers. Additionally, a feature of the present invention is the logging of events on each of the gateway servers, routing servers, and database servers for administrative purposes.

15 In an embodiment, the present invention includes a graphical user interface for configuring the call service properties of a gateway server.

20 In an embodiment, the present invention includes a graphical user interface for listing the country code of a call address maintained by the translation system. The present invention further includes panels for controlling the export of these call addresses as routes. In a further embodiment, the present invention includes a graphical user interface for assigning specific line groups to call address, as well as, controls over the export, timing and access provided to that route.

25 In an embodiment, the present invention includes a graphical user interface for adding, modifying and deleting routes from gateway servers and routing servers.

In an embodiment, the present invention includes a graphical user interface for configuring the features of the gateway server. The present

invention further includes panels for assigning the identity and route import/export features of the gateway server.

In an embodiment, the present invention includes a graphical user interface for listing and/or controlling the handling of local routes.

5 In an embodiment, the present invention includes a graphical user interface for listing and/or controlling the handling of exported local (or dynamic) routes.

In an embodiment, the present invention includes a graphical user interface for listing and/or controlling the handling of static routes.

10 In an embodiment, the present invention includes a graphical user interface for listing and/or controlling the handling of disseminated routes.

Further embodiments, features, and advantages of the present invention, as well as the structure and operation of the various embodiments of the present invention, are described in detail below with reference to the accompanying drawings.

Brief Description of the Figures

The accompanying drawings, which are incorporated herein and form a part of the specification, illustrate the present invention and, together with the description, further serve to explain the principles of the invention and to enable a person skilled in the pertinent art to make and use the invention.

20 In the drawings:

FIGS. 1A and 1B illustrate how a call is generally placed over a VoIP system.

25 FIG. 2 is a diagram of a VoIP system according to one embodiment of the present invention.

FIGs. 3A, 3B, and 3C are diagrams illustrating a gateway server used in the VoIP system of FIG. 2 according to the present invention.

FIG. 3D is a diagram illustrating the routing server used in the VoIP system of FIG. 2 according to the present invention.

FIG. 4 is a diagram illustrating a database server used in the VoIP system of FIG. 2 according to the present invention.

FIG. 5 is an example computer system in one example implementation of the present invention.

FIG. 6A is a diagram illustrating a routine for a gateway server

FIG. 6B is a diagram illustrating a routine for a routing server.

FIG. 7 illustrates a routine for exporting local routes.

FIG. 8 illustrates a routine for transforming exported local routes.

FIG. 9 illustrates a routine for providing disseminated routes.

FIG. 10 illustrates a routine for storing routes.

FIG. 11 illustrates a routine for a routing manager.

FIG. 12 is a diagram illustrating the pre-call processing steps of the present invention.

FIG. 13 is a diagram illustrating the call processing steps of the present invention.

FIG. 14 illustrates a routine for local routing.

FIG. 15 illustrates a routine for static global routing.

FIG. 16 illustrates a routine for disseminated routing.

FIG. 17 illustrates a routine for dynamic routing.

FIG. 18 is a diagram illustrating the E.164 parsing routines of the present invention.

FIG. 19 illustrates a routine for local parsing.

FIG. 20 illustrates a routine for determining a call address.

FIG. 21A illustrates a routine for prioritizing routes.

FIG. 21B is a diagram illustrating the routine of route ordering.

FIG. 22A illustrates the use of multiple gateway servers.

FIG. 22B illustrates the use of a routing server with gateway servers.

FIG. 22C illustrates the use of a database server with gateway servers.

FIG. 23 shows a welcome screen for Dialogic System Software.

FIG. 24 shows a user information screen for Dialogic System Software.

FIG. 25 shows a setup options screen for installation for Dialogic System Software.

FIG. 26 shows an ISDN Protocol Selection screen for Dialogic System Software.

FIG. 27 shows a destination selection screen for Dialogic System Software.

FIG. 28 shows a setup options summary screen for Dialogic System Software.

FIG. 29 illustrates the Dialogic Configuration Manager for Dialogic System Software.

FIG. 30 shows the Protocol Configuration screen for the Dialogic Configuration Manager for Dialogic System Software.

FIG. 31 shows a welcome screen for the driver setup program for Dialogic System Software.

FIG. 32 shows a TAP-H.EXE umcompression screen.

FIG. 33 shows a welcome screen.

FIG. 34 shows an installation destination selection screen.

FIG. 35 shows a welcome screen.

FIG. 36 shows a license agreement confirmation screen.

FIG. 37 shows a welcome screen.

FIG. 38 shows an information screen.

FIG. 39 shows an installation option screen.

FIG. 40 shows a selection screen.

FIG. 41 shows an installation screen.

FIG. 42 shows the express configuration welcome screen.

FIG. 43 shows the gateway server configuration screen.

FIG. 44 shows an identification screen.

FIG. 45 shows a management access screen.

5 FIG. 46 shows a license information screen.

FIG. 47 shows a hardware setup screen.

FIG. 48 shows a call routing screen.

FIG. 49A shows a database access screen.

FIG. 49B shows a client setup screen for configuration of a routing server.

10 FIG. 50 shows a server login setup screen for a database server.

FIG. 51 shows a screenshot of a start menu/programs menu.

FIG. 52 shows a screenshot of the service menu popup menu.

FIG. 53 shows a control panel screen in one embodiment of the present invention.

15 FIG. 54 shows a security domain configuration screen.

FIG. 55 shows a dialog screen.

FIG. 56 shows a remote system list screen.

FIG. 57 shows the service manager screen.

FIG. 58 shows the event logging configuration dialog.

20 FIG. 59 shows a security and identification screen.

FIG. 60 shows a properties screen with identification panel shown.

FIG. 61 shows a properties screen with addresses panel shown.

FIG. 62 shows management access screen.

FIG. 63 shows a client configuration screen.

25 FIG. 64 shows a security identification screen.

FIG. 65 shows a license manager screen.

FIG. 66 shows an address book panel.

FIG. 67 shows a hardware device configuration screen.

FIG. 68 shows a general panel for D/41ESC.

FIG. 69 shows a call analysis panel.

FIG. 70 shows a D/41ESC Ring/Dialtone settings panel.

FIG. 71 shows a tone definition panel.

FIG. 72 shows a cadence dialog.

FIG. 73 shows a D/160SC-LS panel.

FIG. 74 shows a general configuration panel.

FIG. 75 shows a general properties panel for D/240SC-T1.

FIG. 76 shows a channel configuration panel.

FIG. 77 shows the call setup parameters panel.

FIG. 78 displays a MSI/SC board configuration panel.

FIG. 79 shows a voice configuration panel.

FIG. 80 shows a line group window.

FIG. 81 displays a general line group configuration dialog.

FIG. 82 shows a hangup settings panel.

FIG. 83 shows a telephone line resources panel.

FIG. 84 shows a general parameters panel.

FIG. 85 displays a connection detection panel.

FIG. 86 shows an access service window for a management system.

FIG. 87 shows rule panels with settings.

FIG. 88 displays a line group service panel.

FIG. 89 shows an autoforward configuration dialog.

FIG. 90 shows a call rating configuration dialog.

FIG. 91 shows a panel for standard call ratings.

FIG. 92 shows a media manager window.

FIG. 93 shows a new media entry dialog.

FIG. 94 shows a gateway application window.

FIG. 95 shows a country/area code settings panel.

FIG. 96 shows an advanced settings panel for a gateway application.

FIG. 97 shows a user identification window.

FIG. 98 shows configuration screen for user authentication.

FIG. 99 displays a country/area code panel for route translation.

FIG. 100 shows a general panel for country settings configuration.

FIG. 101 shows an address list panel.

5 FIG. 102 displays a table entry for an address panel.

FIG. 103 shows an addressing information screen.

FIG. 104 shows a new route address entry dialog.

FIG. 105 shows an ordering information panel.

FIG. 106 shows a timing information panel.

10 FIG. 107 shows an access control panel.

FIG. 108 displays a route access entry screen.

FIG. 109 shows an export configuration panel.

FIG. 110 shows a configuration screen for a routing manager.

FIG. 111 displays a local route list screen.

15 FIG. 112 shows a routing server settings screen.

FIG. 113 shows an exported routes usage panel.

FIG. 114 shows an exported routes access panel.

FIG. 115 shows an exported routes ordering panel.

FIG. 116 shows an imported routes usage panel.

20 FIG. 117 shows an imported routes access panel.

FIG. 118 shows an imported routes ordering panel.

FIG. 119 shows a routing server configuration window.

FIG. 120 shows a general panel for a routing server.

FIG. 121 shows a routing access panel for a routing server.

25 FIG. 122 shows an exported routes window.

FIG. 123 shows a static route configuration window.

FIG. 124 shows a disseminated route configuration window.

FIG. 125 shows a disseminated routes properties panel.

FIG. 126 shows a database manager configuration screen.

FIG. 127 displays a database server list dialog.
FIG. 128 shows a database server window.
FIG. 129 shows a client configuration dialog.
FIG. 130 shows a server settings window.
5 FIG. 131 shows a user query screen for a provisioning system.
FIG. 132 shows a user identification panel.
FIG. 133 shows a user permissions panel.
FIG. 134 shows a caller ID identification panel.
FIG. 135 shows a caller ID authentication panel.
10 FIG. 136 shows an account record status panel.
FIG. 137 shows a parent group panel.
FIG. 138 shows a group record query screen.
FIG. 139 displays a general panel for group records.
FIG. 140 shows an account record query screen.
15 FIG. 141 displays a general panel for account records.
FIG. 142 shows a profiles and restrictions panel.
FIG. 143 shows a call detail record query screen.
FIG. 144 shows a call detail record service window.
FIG. 145 displays one embodiment for a provisioning system.
20 FIG. 146 displays one embodiment for a management system.
FIG. 147 displays one embodiment for a routing system.
FIG. 148 displays one embodiment for a database server.
FIG. 149 displays one embodiment for a gateway system.
FIG. 150 illustrates a routine for a user account service.
25 FIG. 151 illustrates a routine for a call detail record service.
FIG. 152 illustrates a routine for a event log service.
FIG. 153 illustrates a routine for a multi-database client.
FIG. 154 illustrates a routine for a database access manager.
FIG. 155 illustrates a routine for a service manager.

FIG. 156 illustrates a routine for a database server.

FIG. 157 illustrates a routine for a database client.

FIG. 158 illustrates a routine for a database.

FIG. 159 illustrates a routine for a provisioning system.

FIGs. 160A - B illustrate a routine for a management system.

FIG. 161 displays a block diagram of a provisioning system.

FIG. 162 displays a block diagram of a management system.

FIG. 163 illustrates a routine for a network manager.

FIG. 164 displays a block diagram for a network manager.

The present invention will now be described with reference to the accompanying drawings. In the drawings, like reference numbers indicate identical or functionally similar elements. Additionally, the left-most digit(s) of a reference number identifies the drawing in which the reference number first appears.

Detailed Description of the Preferred Embodiments

Table of Contents

- I. Overview and Discussion
 - A. Open Standards
 - B. Scalability
 - C. Voice quality
 - D. Central or Distributed Implementation
 - E. Built-in Debiting Function
 - F. Multi-Lingual Prompts
 - G. Trunk-to-Trunk Connections
 - H. Universal Connectivity

- II. Terminology
- III. Voice over the Internet System
- IV. Example Implementation
- V. Example Computer System
- 5 VI. Conclusion

I. Overview and Discussion

The present invention provides a method, system, and computer program product for providing management of database servers and services in a voice communication over the Internet system. For convenience, this method, system, and/or computer program product according to the present invention is also referred to herein as simply a "product" or "family of products." The present invention overcomes the above problems in conventional VoIP systems. In one embodiment, there are five basic components:

- Gateway Server,
- Routing Server,
- Database Server,
- Management System, and
- Provisioning System.

These components provide a VoIP system with a modular architecture and a scalable configuration. This modular design creates an efficient environment for fast and effective evolution. New capabilities can be added to the system in a fraction of the normal time. This scalable, platform accommodates future growth in the number of users and other system demands.

In one example, the present invention provides all of the necessary tools to build and operate a large distributed carrier network. The gateway server and database server coordinate to perform calling operations. The routing server

facilitates and coordinates interaction between the gateway servers. The management system is used to manage and configure the other components. Smaller networks can operate with a minimum configuration of only the management and gateway servers - the remaining components are not required to run such a gateway alone.

All five components can run on the same personal computer (PC) or on physically separate machines. For example, a network could typically include a central management server running on the administrator's PC, a number of gateway servers running in various locations and a centralized database server shared between them. The servers/PCs running the various components do not have to be on the same local area network (LAN) as long as they are connected via an IP network.

A. Open Standards

The present invention is compatible in an open architecture. For example, industrial grade IBM-compatible workstations and servers can be used as a hardware platform. Microsoft Windows can be the operating software. Dialogic voice boards that are globally recognized can be used. The availability of Dialogic voice boards for the majority of countries around the world allows an embodiment of the present invention to be distributed on a global basis to customers. Lucent SX-7300P can be used as a CODEC (voice coders and decoders). Its acceptability industry-wide combined with high-quality and low-bandwidth usage makes it a superior choice. Current interoperability voice communication standards: H.323 v2.0 and the voice coders (G.723.1) associated with the standard are also used.

This implementation is not intended to limit the present invention, however, as other types of platforms, operating software, voice boards, CODECs, voice communication standards, and voice coders can be used, as would be

apparent to a person skilled in the art given this description. Proprietary components can also be used instead of or in addition to the open standards described herein.

B. Scalability

5 The present invention further has a scalable architecture. This allows the system to grow with no limitations to thousands of ports regardless of a monolithic or distributed implementation. Growth is modular. In one example implementation, each gateway server module provides up to 4 or more digital boards of T1, T1 PRI, or E1 PRI for easy expansion. This advanced design offers 100% availability as modules are enabled to share the traffic in an unlikely event of a failure of a component. Redundancy of services and back up facilities can be built into the design to provide a high level of reliance for mission critical applications.

C. Voice quality

15 One shortfall of VoIP in the early 1990s was the poor quality of voice (jittery voice) and the unacceptable latency caused by the fluctuating, and at times less than adequate bandwidth available through the Internet.

20 According to a further feature of the present invention, dynamic jitter buffering (DJB) is used to resolve the quality of voice over the unpredictable and some time limited bandwidth of the Internet. This capability dynamically adjusts the size of the voice buffer and eliminates the jitter. In addition, the present invention utilizes the resources of DSP to support verity of voice improvements. A level of quality at or near toll quality is achieved.

D. Central or Distributed Implementation

In one example of the present invention, the VoIP system is mostly implemented in a distributed fashion to enable low cost long distance calling. As a result the organizations using the present invention can form business structures and organizations that are just as innovative and distributed as the VoIP system. While the business needs of customers are different and every changing, they all require systems that are infinitely flexible. For this reason, the present invention in one example is developed with maximum flexibility in mind. The architecture of the system allows for the services such as routing, database, and management servers to reside on one system location or be distributed to several locations. These servers, if connected to the same packet network, will operate as if they are all in the same location.

E. Built in Debiting Function

Many of the early adoptions of Carrier Class VoIP technology were in the area of prepaid long distance. Service providers were often forced to purchase three (3) major components in order to operate their prepaid services. These components were:

- a telephone switch to receive and terminate calls on a trunk to trunk basis;
- a prepaid application server to conduct registry, PIN number generation and validation and account debiting functions; and
- a limited voice over IP gateway to benefit from the cost savings that this technology offers.

The present invention can support all of the above prepaid functions in a gateway architecture and offer them to users (such as customers) at no additional charge.

F. Multi Lingual Prompts

When the present invention is deployed as a network of gateways around the globe, some of the applications and servers can be centralized in one location for ease of operation and savings. Alternatively, the applications and servers may be distributed throughout a network. Either way this allocation of servers is effected by end users which come from many different backgrounds and nationalities. In order to provide the best customer service possible, one embodiment of the present invention uses a full Interactive Voice Response (IVR) system in gateways with multi-language prompts. This enables the VoIP system of the present invention to speak the language of its users and therefore generate more business or traffic.

G. Trunk-to-Trunk Connections

According to a further feature, the VoIP system of the present invention provides VoIP technology even to users (such as businesses) who do not have cost effective and reliable Internet connections available to them at this time but wish to start a prepaid application and move to IP transport later. Prior to the present invention, operators of a start-up VoIP business (e.g., a prepaid service provider) were unable to use any VoIP Gateway to start their prepaid business. The present invention provides PSTN trunk to PSTN trunk connection (no Internet connections) in the same manner as a PSTN switch and offers prepaid services at a fraction of cost. This provides an opportunity for prepaid service providers to establish their business with 80% cost savings.

H. Universal Connectivity

The present invention provides a variety of PSTN and Packet Network connections including, but not limited to, internet, analog PSTN lines, frame relay, PBX extension lines, ATM, digital T1 and T1 PRI, private leased lines, and digital E1 PRI.

The present invention is described in terms of this example environment. Description in these terms is provided for convenience only. It is not intended that the invention be limited to application in these example environments. In fact, after reading the following description, it will become apparent to a person skilled in the relevant art how to implement the invention in alternative environments known now or developed in the future.

II. Terminology

To more clearly delineate the present invention, an effort is made throughout the specification to adhere to the following term definitions as consistently as possible.

The terms "packet-switched network" and "Internet" are used interchangeably to refer to any network that performs packet switching including, but not limited to, a network that supports a Transmission Control Protocol (TCP) and/or an Internet Protocol (IP). The terms "packet-switched network" and "Internet" are not intended to limit the size of a network, and can include, but are not limited, to a wide area network (WAN), middle-area or campus wide network, or local area network (LAN), or combinations of networks.

The term "traffic" refers to voice, facsimile, video, multimedia, digital information, or other data that can be sent between telephony terminal equipment and/or network terminal equipment.

The term "routes" refers to any listing of address information, either network or telephony or both, which can be used to direct traffic.

The term "window" and "screen" are used interchangeably and refer to a graphical user interface for expressing, confirming, and/or altering the configuration and/or operation of the present invention.

The term "panel" refers to any graphical user interface component displayed within a window or screen.

The term "dialog" refers to a window or screen which is displayed while interacting with a window or screen in any graphical user interface such that configuration can occur. Dialogs are usually displayed with Ok and Cancel buttons. Dialogs usually return a user to the original window or screen after confirmation or negation.

The term "environment" collectively refers to the components of the system and the connections among them. "Software environment" refers to the environment constructed by the various processes and routines of the system.

The use of the terms "view" and "document" refer to the objects used in computer programming. A document is an object for storing data. A view is an object for rendering visual representations of a document's data on a display, either graphically, textually, on screen, in print, or the like. Views are also used to translate user input into the system. For example, by recording mouse movements and keyboard messages. A detailed treatment of the properties of these objects can be found in *Programming Windows 95 with MFC, Part VII: The Document/View Architecture* by Jeff Prosise, Microsoft Systems Journal, 1995.

III. Voice over the Internet System

FIG. 2 is a diagram of a VoIP system 200 according to an embodiment of the present invention. VoIP system 200, among other things, allows traffic (e.g., voice or fax data) originating on a circuit-switched network to be carried over a

packet-switched network. In particular, VoIP system 200 acts as a bridge between a public switch telephone network 205 and an IP network 215

Such voice and data traffic can originate at any type of terminal equipment at PSTN 205. For example, a call can be placed on an ordinary telephone 201 or 202. A call can be placed from an ordinary telephone 203 using a prepaid service and autodialer 204. Calls can also originate from a speaker or conference-type phone 206 through a private branch exchange (PBX) 207. Such calls can terminate at an IP phone 293 or an ordinary telephone 291, 292. These examples are not intended to limit the present invention. As will be apparent to a person skilled in the art given this description, any type of computer or telephone equipment can be used to originate and terminate voice calls over the VoIP system 200.

VoIP system 200 includes Gateway Servers 210, 220, Database Server(s) 240, Routing Server(s) 230, Management System(s) 250, Provisioning System(s) 260, Network Manager(s) 270, and License Server(s) 280. Each of these components can be interconnected over the IP network 215, as shown in FIG. 2. In one embodiment, Provisioning System 260 is coupled to Database Server 240. Provisioning System 260 is controlled by graphical user interface (GUI) tools 265. Similarly, Management System 250 is controlled by graphical user interface tools 255. A relational database such as a SQL Database 245 is further coupled to Database Server 240.

Another embodiment employs only one Gateway Server 210 with the above components in the same configurations. The lack of second Gateway Server 220 would restrict the economy of the two or more gateway server embodiment, since long distance charges are more likely to occur when the same gateway handle incoming and outgoing telephone calls. However, a single gateway server would suffice for a working VoIP system, and would be feasible in some situations. For instance, call center applications, which entail the use and tracking of internal telephone networks are one possible implementation of this

single gateway embodiment. Others are possible and this implementation is merely an example.

VoIP system 200 is designed to allow voice and fax traffic to be carried over IP network 215. Voice and fax traffic received from PSTN 205 is coupled through gateway servers 210 and 220 over IP network 215. In addition to providing simple call termination, VoIP system 200 provides a host of enhanced features that enable IP telephony capabilities to be used in a number of different applications. For example, since IP telephony typically results in much lower telecommunications costs, and can be used between two branch offices to reduce costs. This reduced cost can instead be used to resell long-distance service to end-users. Another application is to enable IP telephony terminals, such as multimedia computers running group meeting software, to access traditional services, such as a call center. By combining IP telephony with other technologies, such as the World Wide Web (WWW), integrated WWW-VoIP system 200 makes these applications possible by acting as a "gateway" between traditional telecommunication networks (such as PSTN 205) and systems, and a packet switch network (such as IP Network 215) and systems.

In one example implementation, there are two main parts to VoIP system 200: a software part and a hardware platform.

An overview of each of the components of VoIP system 200 is now provided. Further descriptions with respect to an example implementation is provided below in Section IV. Gateway Server 210 (and Gateway Server 220) are coupled between PSTN 205 and IP Network 215. Gateway Server 210 can be coupled to any number of Telephone Lines 362 and Network Lines 364.

Gateway Server 210 includes, in one embodiment, both software and hardware components. For example, as shown in FIGs. 3A-3C, Gateway Server 210 includes a Gateway Server Application Layer 302, Gateway Server Platform 304 (software), and Gateway Server Hardware 306. Gateway Application Layer 302 includes Higher Level Application Services 310. For example, these

application layer services can include Gateway Service 311, Common Services 312, Auto-Forward Services 313, Other Services 314, Higher Level Services 315, and Abstract Services 316. In one example implementation where the software is written in an object-oriented program language such as C++, each of the Higher Application Layer Services 310 accesses lower-level modules (also called objects) through a Session Manager Application Programming Interface (API) 320.

Gateway Server Platform 304 includes Session Manager 330. Session Manager 330 includes Rule Engine 332. Session Manager 330 creates and manages Sessions 340, 341, 342. In general, any number of sessions can be managed, depending on the resources of the Gateway Server 210. Each session corresponds to a voice call carried over Gateway Server 210. Information in a Session 340, 341 or 342 can be further handled by lower level Line Group Manager 350, Routing Manager 352, Database Access Manager 354, Media Manager 356, and Call Rating Manager 358. Routing Manager 352 is further coupled to a Subsystem 351. Subsystem 351 can contain various components dealing with the implementation and execution of routes. For instance, Local Parsing 1170 (FIG. 19) and Wildcarding 2002 (FIG. 20) can be included in Subsystem 351. FIGs. 19 and 20 are discussed in the E.164 Parsing section below. Database Access Manager 354 is also coupled to a Subsystem 353. For example, Subsystem 351 can be any cache memory or other storage device for storing routing information. Subsystem 353 can be a relational or nonrelational database.

Line Group Manager 350 coordinates communication between the telephone line side and network line side of Gateway Server 210 for a particular session or call. In this way, the Higher Application Layer Services 310 and Session Manager API 320 can be transparent to any particular protocols required by a telephone line or a network line.

Conversion Module 367 then converts a telephone line group to a line group. Conversion Module 368 converts a network line group to a line group. In this way, Line Group Manager 350 manages line groups without regard to whether a line group is a particular telephone line or network line. That is, the type of line is transparent to the Line Group Manager 350 and any other higher-level application services.

Hardware Device Manager 360 is the module or object responsible for managing the connectivity of Gateway Server 210 to traditional telecommunication devices and to conventional internet-related networking components. Thus, Hardware Device Manager 360 coordinates with different respective modules or drivers to control telephony and network components. For example, as shown in FIG. 3C, Gateway Server Hardware 306 can include a NIC card 381 as part of PC. NIC Card 381 is further supported by Gateway Server Software 304 components to coordinate information sent to Hardware Device Manager 360. In particular, a TCP/IP Stack 376, NetSupport Module 377, H 323 Stack 374, and H.323 Support Module 375, are used to couple NIC Card 381 and Hardware Device Manager 360. Similarly, for telephony components, Gateway Server 210 includes a Dialogic telephony board(s) 382 that can be coupled to phone lines over PSTN 205. Gateway Server Hardware 306 can further include an analogic Brooktrout Digital Signal Processor (DSP) 383. Dialogic Telephony Board(s) 382 are supported by respective Dialogic Drivers 370 and Dialogic Middleware 371. Similarly, Analogic Brooktrout DSP(s) 383 are supported by a DSP Driver 372 and a respective DSP Middleware Support Module 373. This example including the specific vendor products are illustrative. Other software and hardware modules supporting internet and telephony services can be provided in Gateway Server 210, as will be apparent to a person skilled in the art given this description.

The embodiments of the present invention are discussed in detail below with reference to the following figures. FIGS. 3A - 3C, 6A, 11, 19, 20, 21A,

21B, 22A, and 149 are described with respect to the gateway server and its features and components. FIGS. 3D, 6B, 7 - 10, 22B, and 147 are described with respect to the routing server and its features and components. FIGS. 4, 22C, 148, and 150 - 158 are described with respect to the database server and its features and components. FIGS. 12 - 18 are described with respect to the gateway server, routing server and their embodiments. FIGS. 145, 159, and 161 are described with respect to the provisioning system and its features and components. FIGS. 146, 160A, 160B, and 162 are described with respect to the management system and its features and components. FIGS. 163 and 164 are described with respect to the network manager and its features and components.

Furthermore, FIGS. 23 - 41 are described with respect to the installation of computer software components of the present invention. FIGS. 42 - 66 are described with respect to the configuration of the features and components of the gateway server, routing server, and database server. FIGS. 67 - 79 are described with respect to the configuration of driver resources. FIGS. 80 - 88 are described with respect to the configuration of line groups. FIGS. 89 - 93 are described with respect to the service modules. FIGS. 94 - 118 are described with respect to the configuration of features and properties of the gateway server. FIGS. 119 - 125 are described with respect to the configuration of routes. FIGS. 127 - 144 are described with respect to the configuration of the database server and its properties and connections to the other components of the VoIP system.

The above listed figures are now be described in detail by subject.

FIG. 149 is a block diagram of another embodiment of the gateway server. It includes Gateway Server Application Layer 302 and Gateway Server Platform 304. In this embodiment, Hardware Support 14936 is shown for Line Group Manager 350. Hardware Support 14936 provides the necessary functionality for the gateway to interact with the network and encompasses Line Group 366, 368, Lines 362, 364, and Hardware Device Manager 360. E.164 Parser 14956 provides routing capabilities to the gateway server. E.164 Parser 14956 performs

local parsing routine 1902, as shown in FIG. 19, and discussed above with respect to Subsystem 351. Internal Services Manager 14820 coordinates among the components of the operational systems: gateway and database. In FIG. 149, Internal Services Manager 14820 coordinates the access of Services 311, 312, and 313 with SM API 320. Other services exist and are shown in FIG. 3B. Internal Services Manager 14820 manages these application layer services interaction with the gateway servers and database servers as well.

The operation of Gateway Server 210 is described further below with respect to an example implementation.

FIG. 3D is a block diagram of the Routing Server 230. Routing Server Hardware 394 can be the same hardware as Gateway Server Hardware 306, however Computer System 500 is sufficient. Routing Server Platform 392 comprises program interfaces and modules for the storage of routing information. This routing information is used by Routing Server Application Layer 390 for storing, serving, transforming and parsing both network and telephony routes.

FIG. 4 is a diagram of an example Database Server 240 according to the present invention. Database Server 240 includes database access manager 354. Database Access Manager 354 further coordinates and operates higher-level services including: a Call Detail Reporting (CDR) Service 422, Event Service 424, and User Service 426. Each of these services further coordinate with a Multi-Database Client 430. Multi-Database Client 430 is coupled to three Database Clients 442, 444, and 446. Each of these Database Clients 442, 444, 446 is further coupled to communicate with respective Database Client Servers 452, 454, 456. Database Clients 442, 444, 446 can be software or hardware applications running locally or remotely. Database Clients 442, 444, 446 can provide access to Database Client Servers 452, 454, 456 through a firewall. Database Client Server 452, 454, 456 are database servers configured to connect only to a specific database clients. Database Client Server 452, for example, can access a Relational Database (SQL) 462. Database Client Servers 454 and 456

likewise can access respective relational or non-relational databases (not shown). Multi-Database Client 430 is not limited to three Database Clients 442, 444, and 446. It can connect to any number of database clients similar to 442, 444, and 446. Similarly, any number of Multi-database Clients 430 can operate together from various database servers 240.

FIG. 148 is a block diagram which shows in another embodiment further detail of the components of Database Client Servers 452, 454, and 456. Database Service 14804 begins with the instantiation of Database Manager 14812. Once started, Database Manager 14812 contacts Software Environment 14516 and Thread Pool Queue 14808 which activates the database servers within the system. Software Environment 14516 is the rest of the VoIP system 200. Thread pool queue 14808 feeds requests to request handler 14816. Request Handler 14816 forward requests to service modules 14832 and incidences of service requests to internal services manager 14820. Service modules 14832 include call detail record (CDR) service 422, event log service 424, user account service 426, and/or other services 314 which need to be implemented in order to handle a request. Requests are forwarded to database manager 14812 by request handler 14816 from service modules 14832.

Database connection manager 14828 and database connection(s) 14848 provide service modules 14832 with the appropriate format for sending and receiving requests from database clients 442, 444, and 446. Database connection manager 14828 and database connection 14848 retrieve the appropriate formats from Open Database Connectivity (ODBC) API 14852.

FIG. 150 illustrates a routine for User Account Service 426 (routine 15002). Database Server 240 receives a user request from a gateway server (step 15004) and obtains user information provided to the gateway server by the user (step 15006). Information, e.g., user information, to fulfill the request is provided by the Multi-Database Client 430 (step 15008). Multi-Database Client 430 forwards the request for information to Database 462. The request is forwarded

through Database Client 442 and Database Client Server 452 to Database 462. Database Clients 444, 446 and Database Client Servers 454, 456 work similarly. Requested information is returned to the service level by Multi-Database Client 430 (step 15010) and then back to the gateway server from which the request originated for confirmation (step 15012).

FIG. 151 illustrates a routine for Call Detail Record (CDR) Service 422 (routine 15102). Gateway Server 210 receives a call and Internal Services Manager 14820 initiates the recording of the call's details by CDR Service 422 (steps 15104 and 15106). CDR Service 422 stores CDRs on CDR Server 423 (step 15108) and when available sends the CDRs to Multi-Database Client 430 (step 15110). Stored CDRs are cached until Database Access Manager 354 requests (step 15112) that they be forwarded to Multi-Database Client 430 (step 15114).

FIG. 152 illustrates a routine for Event Log Service 424 (routine 15202). Gateway Server 210 receives a request for service and Event Log Service 424 is notified to record a log of the activities pertaining to the request (step 15204). Event Log Service 424 determines the type of service being requested, e.g., user account and call detail record are just two types of services, (step 15206) and records the request (step 15208). Event Log Service 424 sends the recorded information, in the form of a log record, to Multi-Database Client 430 for storage in Database 462 (step 15210).

FIG. 153 illustrates a routine for Multi-Database Client 430 (routine 15302). Multi-Database Client 430 connects to the available service module(s) (step 15304) and the available database client(s) (step 15306). Multi-Database Client 430 receives requests from the service module(s) (step 15308) and queues the requests (step 15310). Multi-Database Client 430 forwards the requests to the appropriate database client(s) (step 15312) and maintains a link until it receives confirmation from the database client(s) that the database has been reached and the information determined (step 15314).

FIG. 154 illustrates a routine for Database Access Manager 354 (routine 15402). Database Access Manager 354 connects to service modules (step 15404) and monitors them (step 15406). When a service module needs access to a component of the VoIP system, Database Access Manager 354 authorizes the service modules access if so configured (step 15408). Database Access Manager 354 also prioritizes requests from service modules and monitors their continued access (step 15410).

FIG. 155 illustrates a routine for Internal Service Manager 14820 (routine 15502). Internal Service Manager 14820 connects to Request Handler 14816 (step 15504), Service Modules 14832 (step 15506), Multi-Database Client 430 (step 15508) and Database Access Manager 354 (step 15510). Internal Service Manager 14820 forwards requests for service from Service Modules 14832 (step 15512), and confirms the handling of the request from Request Handler 14816 (step 15514).

FIG. 156 illustrates a routine for Database Client Server 452 (routine 15602). Database Client Server 452 begins with the execution of Database Service 14804 (step 15604). Database Service 14804 instantiates Database Manager 14812 (step 15606). Database Manager 14812 builds Thread Pool Queue 14808 (step 15608). Both Database Service 14804 and Database Manager 14812 connect to Database 462 through Software Environment 14516 (step 15610). In step 15612, Database Manager 14812 connects with Database Client 442. Database Manager 14812 maintains Thread Pool Queue 14808 (step 15614) and confirms the handling of requests from Request Handler 14816 (step 15616). Database Manager 14812 receives requests from Database Client 442 (step 15618) and after accessing Thread Pool Queue 14808 to reach Database Connection 14848, receives a response from Request Handler 14816 and is able to respond to Database Client 442 (step 15620). Additional requests are forwarded to Thread Pool Queue 14808 (step 15622) or responded to Database

462 through Software Environment 14516 (step 15624). From either step 15620 or 15624, responses are forwarded to Multi-Database Client 430 (step 15626).

FIG. 157 illustrates a routine for Database Client 442 (routine 15702). Database Client 442 connects to Database Service 14804 (step 15704) and is given access to Database Manager 14812. Database Client 442 also connects to Multi-Database Client 430 (step 15706) and Software Environment 14516 (step 15708). Database Client 442 receives requests from Multi-Database Client 430 (step 15710) and Software Environment 14516 (step 15712). These requests are forwarded to Database Client Server 452 (step 15714) and responses given from Database Client Server 452 (step 15716). Database Client 442 forwards the response to Multi-Database Client 430 (steps 15718) and Software Environment 14516 (step 15720).

FIG. 158 illustrates a routine for Database 462 (routine 15802). Database 462 connects to Database Client Server 452 (step 15804). Database 462 receives requests from Database Client Server 452 (step 15806) and searches for the requested information (step 15808). In some circumstances, Database 462 creates a new record (step 15810) or modifies an existing record (step 15812). In step 15814, Database 462 responds to Database Client Server 452 with the information that was requested in step 15806.

As shown in FIG. 2, Gateway Server 210 manages the specialized telephony hardware that is used for interconnection to telecommunications equipment, and performs bridging of calls between the traditional telecommunications network and an IP network to which Gateway Server 210 is connected. The key responsibilities of Gateway Server 210 are: handling calls received on either IP network 215 or PSTN 205, bridging calls between the two different kinds of networks (circuit switched and packet switched). Such bridging can include, but is not limited to, voice compression, echo cancellation, control protocol transformation, and numerous other complex operations, interacting with telephone users using DTMF digits (and other methods of telephony signaling)

in order to collect information such as a target destination number, user identification, password and more.

Gateway Server 210 is also responsible for interfacing with a routing system (described below) in order to translate numbers and routing those numbers to the correct destination gateway, IP address, or telephone number, depending on the number dialed. Gateway Server 210 interfaces with Database Server 240 in order to validate user accounts, apply charges, and keep records of all calls.

Routing Server 230 is used to route calls across IP network 215, and possibly to a final destination on the telephone network. Routing Server 230 is optional in that Gateway Server 210 can route a call based on local information maintained on Gateway Server 210. However, Routing Server 230, according to a further feature of the invention, can facilitate the automatic exchange of routes between different gateway servers 210, 220. This in general can make the administration configuration of call routing much easier.

FIG. 147 is a block diagram of an embodiment of the routing system. Routing services 14704 instantiate Routing Server 230 and links to software environment 14516. Routing Server 230 also links to software environment 14516 and communicates with Routing Client(s) 14712. In one embodiment, the Clients of the Routing Servers Are Gateway Server(s) 210, 200. Routing Server 230 has access to the same session manager API 320 that is used in Gateway Server 210. Thus, routing server(s) 230 can perform similar operations.

In this way, routing server(s) 230 can serve the following functions: They can look up routes to a particular number; they can allow Gateway Server 210, 220 to export local routes for the use of other gateways in the system; and Routing Server 230 can further provide Gateway Servers 210, 220 with routes that are specifically intended for that respective gateway server. These functions allow all routes to be configured on the Routing Server 230 only, but automatically distributed to the appropriate gateways. Routing Server 230 can also distribute E.164 translation data through a E.164 Parser in Subsystem 351.

This facilitates the exchange of such data, eliminating the need to configure it on each Gateway Server 210, 220 individually.

According to a further feature, in an enlarged deployment, and for reliability, it is possible to configure Multiple Routing Servers 230 in the same network. Multiple servers can provide both load-balancing across multiple machines and redundancy to protect against the failure of any particular Routing Server 230.

Thus, Routing Server 230 is responsible managing routing clients. In one embodiment, Gateway Server(s) 210 are the routing clients. Routing Server 230 provides new and updated routes to the routing clients to which it is assigned. Routing Server 230 also has the ability to perform searches for a matching routes within its own stored directory of routes. Routing clients can request such searches of Routing Server(s) 230.

Database Server 240 manages the users of VoIP system 200. Database Server 240 supports the restriction of calling privileges to users with valid user accounts. In addition, Database Server 240 allows billing for VoIP services to be done in either a postpaid or in a real-time manner. Alternatively, Gateway Servers 210, 220 can be connected directly to a database, such as Database 245. Preferably, however, Database Server 240 is provided so that VoIP system 200 can be scalable. As described with respect to FIG. 4, Database Server 240 can scale to accommodate any number of databases depending on the number of users and/or services provided by VoIP system 200.

In one example, Database Server 240 provides Multi-database Client 420 as middleware for interfacing with numerous Gateway Servers 210, 220 while using only a small number of connections to the underlying database(s). This allows even very large numbers of Gateway Servers 210, 220 to have access to data contained in a shared database or databases. Multiple Database Servers 240 can be used as shown to provide redundancy and protection against failure.

Management System 250 uses GUI Tools 255 to enable configuration administration of software components for the VoIP system 200. Management System 250 can connect to gateway servers 210, 220, Routing Server 230, and Database Server(s) 240 remotely over IP Network 215. In this way, Management System 250 can monitor the current state of these servers and make changes to the configuration of these components. This architecture allows all of the Gateway Servers 210, 220, Routing Server 230 and Database Server 240 in the network to be configured from a single point.

FIG. 146 is a block diagram showing one embodiment of Management System 250. Management System Document 14612 links to Software Environment 14516 to obtain information (data) from the various components of the VoIP System 200. In one embodiment, the Management System Document 14612 links to Software Environment 14516 and to Management System View 14604. Management System View 14604 provides the interface from which windows, screens and displays can be presented. These windows, screens and displays are presented to users.

Users can interact with VoIP system 200 with the GUI Tools 255 presented through Management System View 14604. Management System view 14604 provides users with the screens, windows and displays to modify the VoIP System 200 and forwards configuration information to Management Client 14620. Management Client 14620 interacts with Software Environment 14516. In one embodiment, Management Client 14620 would interact with the Database Server 240 which is a component of Software Environment 14516. If needed, Management Client 14620 would relay information to Management Client Object 14616 and updates to ActiveX/OCXs 14624.

ActiveX controls (formerly OLE controls) are objects that can be inserted into any application that is an ActiveX control container like IE, Visual Basic and Delphi. Examples include buttons, stock tickers, and chart controls. For example, Animation GIF ActiveX Control exposes the AniGIF object that you

can use to create new applications and programming tools with Animation GIF interface. Many development platforms now support ActiveX Controls. These containers include Visual Basic, Visual C++, Visual Basic for Application 7.0, Visual Foxpro, Access, Delphi, PowerBuilder, and many others. OCXs are Object Linking and Embedding Controls and are known commonly as .ocx files.

ActiveX/OCXs 14624 are stored, updated as shown, and implemented in management client objects 14616 and management controls 14608. Management controls 14608 receives implements from ActiveX/OCXs 14624 and events from management client objects 14616. Management controls 14608 manages the configuration of the VoIP system so that conflicts among system components do not occur.

Provisioning System 260 is a tool that allows access to the data managed by Database Server 240. Management System 250 provides access to the Database Server 240, and allows Database Server 240 to be configured, but it is the role of Provisioning System 260 to actually allow interaction with data contained in Database(s) 245. In one embodiment, Provisioning System 260 performs the following tasks using GUI Tools 265: Provisioning System 260 allows accounts of system users to be added, removed or modified; Provisioning System 260 enables a user to configure and maintain group records of users, and account records of users; and Provisioning System 260 also enables call detail records stored by the Gateway Servers 210, 220 to be reviewed. Like Management System 250, Provisioning System 260 operates remotely over IP network 215, and can be used to administer multiple Database Servers 240.

FIG. 145 is a block diagram showing one embodiment of Provisioning System 260. Provisioning System Document 14508 links to software environment 14516 to obtain information (data) from the various components of the VoIP system 200. In one embodiment, the Provisioning System Document 14508 links to Software Environment 14516 and to Provisioning System View 14504. Provisioning System View 14504 provides the interface from which

windows, screens and displays can be presented. These windows, screens and displays are presented to users.

Users can interact with VoIP System 200 through Provisioning System View 14504. Provisioning System View 146504 provides users with the screens, windows and displays to modify the VoIP System 200. Database Client Subsystem 14512 is linked to Software Environment 14516. Database Client Subsystem provides data from the database server to requesting components of the VoIP System 200. Users can access and view this data from Provisioning System view 14504.

Network Manager 270 is a separate application that allows VoIP System 200 to be configured or modified in response to specific client requests. FIG. 164 shows a block diagram of one embodiment of Network Manager 270. A client request made in the form of a change at a Client Database 275 is automatically recognized by Network Manager 270. Network Manager 270 then formulates and propagates appropriate control signals to Software Environment 14516 in order to configure VoIP system 200 accordingly. In this way, specific client requests and needs are automatically made to VoIP system 200.

FIG. 163 illustrates a routine for Network Manager 270. Network Manager 270 queries Client Database 275 (step 16304) and receives data from Client Database 275 (step 16306). Network Manager 270 determines from this data if an update is needed (step 16308). If not, then Network Manager 270 returns to querying (step 16310). If so, then Network Manager 270 sends a message representative of the update to Software Environment 14516 (step 16314). Network Manager 270 further sends the message to available gateway server(s) (step 16316), database server(s) (step 16318), and routing server(s) (step 16320).

Licensing Server 280, according to a further feature of the present invention, allows licensing of the present invention technology according to license keys. A license key is an encoded string that contains information

regarding a particular license. License Server 280 is an essentialized license server that manages licensing for all components in VoIP System 200. When a component first starts up, it will contact License Server 280 to receive necessary license. With this configuration, all license keys are configured on a single machine, and only one hardware key is required.

IV. Example Implementation

Table of Contents

I. BASICS OF THE PRESENT INVENTION

1. Introduction
2. System Overview
 1. The Present Invention
 2. Present Invention Components
 - a. Gateway Server
 - b. Routing Server
 - c. Database Server
 - d. Management System
 - e. Provisioning System
 - f. Utility Components
 3. Other Implementation Details
 - a. Components and Nodes
 - b. Operational vs. Administrative Components
 - c. Configuration
 - d. Identification and Authentication
 - e. Access Control
 - f. Management
 - g. Licensing
3. Installation and Setup

1. Hardware Installation and Setup
 - a. Installing Dialogic Telephony Boards
 - b. Installing DSP Boards
 - c. Installing the SCbus cable
 - d. Installing the Dialogic Drivers
 - e. Installing Analogic Drivers
 - f. Installing the Brooktrout Drivers
2. Software Installation
3. Express Configuration
 - a. Running Express Configuration
 - b. Common Setup Information
 - (1) Identification
 - (2) Management Access
 - (3) License Information
 - c. Gateway Server Setup
 - (1) Hardware Setup
 - (2) Cell Routing
 - (3) Database Access
 - d. Routing Server
 - e. Database Server
 - (1) Client Setup
 - (2) Server Login Setup
4. Database Table Setup
5. Additional Setup Issues
 - a. Screen Savers
 - b. Find Fast / Virus Scanners
4. Getting Up and Running
 1. Starting the Software
 2. The Array Control Panel

- 3. Using the Management System
 - a. Initial Configuration
 - b. Remote System List
 - c. Connecting and Administering
- 5 4. Service Control and Logging
 - a. Service State Control
- 5. Security Domain Configuration
 - a. Memberlist Domains
- 6. Management Access
- 10 7. Licensing
 - a. Standalone Mode
 - b. Client/Server Mode
- 8. Using the Provisioning System
- 15 9. Troubleshooting Service Startup
 - a. Gateway Service
 - b. Routing Service
 - c. Database Service
- II. OPERATION
- 20 1. Gateway Server Capabilities
 - 1. Hardware Support
 - a. Gatenet 3 vs. DSP Compression
 - b. Dialogic D/21H
 - c. Dialogic D/41ESC
 - d. D/41ESC General Panel
 - 25 e. D/41ESC Cell Analysis Panel
 - f. D/41ESC Ring/Dialtone Settings Panel
 - g. D/41ESC Tone Definition Panel
 - h. Dialogic D/160SC-LS
 - i. Dialogic D/240SC-T1 (Robbed Bit T-1)

- j. Dialogic ISDN
 - (1) General Panel
 - (2) Channel Configuration Panel
 - (3) Cell Setup Parameters Panel
- k. Dialogic MSI/SC
- l. Mibridge DSP
- m. Network H.323
- n. Network Proprietary
- 2. Lines and Line Groups
 - a. Lines
 - b. Line Groups
 - c. Telephony Lines and Line Groups
 - (1) Hangup Settings Panel
 - (2) Integration Digital Settings
 - (3) Resources Panel
 - (4) General Panel
 - d. Cell Analysis
 - (1) Connection Detection Panel
 - (2) Disconnection Detection Panel
 - (3) Network Lines and Line Groups
- 3. Services
 - a. Service Selection
 - b. Service Level Line Group Settings
 - c. Gateway Service
 - d. Autoforward Service
 - e. Internal Services
- 4. Call Rating
- 5. Voice Prompt Management
- 2. The Gateway Application

1. Principles of Operation
2. Gateway Application Configuration
 - a. General Settings Panel
 - b. Country/Area Code Settings Panel
 - c. Advanced Panel
3. Supplementary Application Configuration
 - a. User Identification Application
 - b. User Authentication Application
4. Session Bindings
5. Routing and Number Translation
 - a. Number Translation
 - b. E.164 Parser
 - c. Prioritization and Examples
 - d. Alternative Configurations
 - e. E.164 Parser Data Synchronization
6. Routes and Route Information
 - a. Addressing Information
 - b. Ordering Information
 - c. Timing Information
 - d. Access Control Information
7. Routing Models
 - a. Locally Defined Routes
 - b. Static Global Routes
 - c. Disseminated Routes
 - d. Dynamic Routes
 - e. Dynamic Caching
 - f. Summary
8. Routing Manager
 - a. Local Route Configuration

	b.	Dynamic Cached Routes
	c.	Routing Server Settings
5	(1)	General Panel
	(2)	Exported Routes Usage
	(3)	Exported Routes Access
	(4)	Exported Routes Ordering
	(5)	Imported Routes Usage
	(6)	Imported Routes Access
	(7)	Imported Routes Ordering
10	9.	Routing Servers
	a.	Routing Clients
	(1)	General Panel
	(2)	Routing Usage Panel
	(3)	Routing Access Panel
15	(4)	Exported Routes Ordering
	b.	Routing Groups
	c.	Exported Routes
	d.	Static Routes
	e.	Disseminated Routes
20	(1)	Route Management Embodiment
	3.	Database Basics
	1.	Gateway Access to Database Servers
	2.	Database Server General Setup
	3.	Database Server Access Control
25	4.	User Management
	1.	User Management Concepts
	a.	User Records
	b.	Group Records
	c.	Account Records

- d. Remote Gateways and User Management
- 2. Gateway Server Access to Users
- 3. Managing User Records
 - a. User Record Queries
 - b. User Record Information
 - (1) Identification Panel
 - (2) Permissions Panel
 - (3) Cellular ID Identification Panel
 - (4) Cellular ID Authentication Panel
 - (5) Account Record Status Panel
 - (6) Parent Group Panel
- 4. Managing Group Records
 - a. Group Record Queries
 - b. Group Record Information
 - (1) General Panel
- 5. Managing Account Records
 - a. Account Record Queries
 - b. Account Record Information
 - (1) General Panel
 - (2) Profiles and Restrictions Panel
- 5. CDR Collection
 - 1. Querying CDRs on a Database Server
 - 2. Information in a CDR
 - 3. Gateway Server Collection of CDRs
 - 4. Database Server Failure Handling

6. License Server

III. APPENDIX A

1. Database Table Descriptions

1. ACCOUNT_RECORD Table

- a. Table Definition
- b. Additional Notes

2. AUTHENTICATE_ID Table

- a. Table Definition
- b. Additional Notes

3. CALL_RECORD Table

- a. Table Definition
- b. Additional Notes

4. CREDIT_RECORD Table

5. DEBIT_RECORD Table

- a. Table Definition
- b. Additional Notes

6. EVENT_LOG Table

- a. Table Definition

7. GROUP_RECORD Table

- a. Table Definition

8. USER_RECORD Table

- a. Table Definition

IV. Appendix B

1. Network Design Information

1. Network Requirements

- a. Gateway Server
- b. Routing Server
- c. Database Server

- d. Management System
- e. Provisioning System
- 2. Traffic Prioritization

I. BASICS OF THE PRESENT INVENTION

1. Introduction

This section provides an introduction to the present invention.

There are three main parts as follows.

Part 1 – Basics of the present invention, provides basic information that applies to all of the components of the present invention. This part contains three sections:

System Overview – Provides an overview of the present invention, describing the key components of the present invention, and how these components interact to provide flexible IP Telephony services. This section also discusses common concepts that are helpful in order to understand the interaction between present invention components. These concepts include identification, security, configuration management, services, and management.

Installation and Setup – Describes how to install and configure the present invention. This section discusses hardware board installation, hardware driver installation and setup, as well as the installation and basic configuration of the present invention software. Only basic setup is covered in this section; later sections provide full details on how to configure the present invention .

Getting up and Running – Completes the introduction to basic functionality by demonstrating how to get a basic system up and running after it has been installed and set up. Shows how Management System 250 is used to

configure operational components, how services can be started and stopped, and how common functionality (such as identification and security) can be set up.

Part 2 – Operation discusses how to use the functionality provided by the present invention in detail. There are four sections in this section, each discussing a specific aspect of functionality provided by VoIP System 200. The sections are as follows:

Gateway Server Capabilities – Provides a detailed description of the gateway server capabilities. This section discusses how to control the interaction of the end users with the system, and how to configure items like voice prompts, services, and numerous others.

The Gateway Application – Introduces the main application that runs on the gateway server, the gateway application, which provides the main service offered by the present invention. Discusses the configuration and operation of the gateway application in detail.

Routing and Number Translation – The present invention provides numerous ways to translate digits dialed by the user and to route calls across a network of gateways. This section describes the routing process, and explains how this process can be optimally configured for a particular environment and network structure.

Database Basics – Provides an introduction to the database servers, and describes how the gateway servers makes use of them. This section introduces some of the basic information, as a prerequisite to discussion features of which the database server is an integral part, in the following sections.

User Management – An important requirement, particularly for telecommunications service providers, is the management of the end users of the system, and the maintenance of their accounts. This section discusses user management, explains the various information maintained for each user of the system, and shows how this information can be used to customize the service provided to the end user.

CDR Collection – In order to bill for, and have a record of calls, the present invention is capable of collecting a Call Detail Record (CDR) for each call. This section explains the CDR collection process, and how it can be customized for a particular environment.

5 *License Server* – In some environments, a license server can facilitate the network-wide licensing of present invention software. This section provides details on how to set up a license server for the purpose of simplifying license administration.

10 Part 3 – Technical Information provides detailed technical information about the present invention. This information is useful for optimizing the performance or quality of an present invention installation, or for integrating external applications with the present invention. There are two appendices in this section:

15 *Database Table Descriptions* – The structure of database tables used by the database server is described in Appendix A. This information is useful in developing customized applications that interface directly with the data stored in the database.

20 *Designing a Network* – The quality of any IP Telephony system is heavily dependent on the quality of the underlying IP network used to transport the IP traffic. This section provides a detailed description of the network factors that affect the overall performance of the system, and provides information on bandwidth and port numbers necessary to design a network that will deliver maximal quality.

2. *System Overview*

25 This section provides an overview of an example implementation of the present invention. It describes the components that make up the present invention, and how these components co-operate in order to deliver IP Telephony

services. It describes the basic concepts that are common to all components of the present invention.

There are three main sections in this section. They are as follows:

5 **The Present Invention.** This section describes, at a very high level, what the present invention is, and what it provides the capability to do.

Present Invention Components. This section introduces the components of the present invention. It describes the purpose of each of the components, as well as the interaction between those components.

10 **Key Concepts.** The key concepts that are common to all components of the present invention are described in this section. The concepts covered include identification, security, configuration, management, and licensing. The basic principles of operation, such as starting and stopping services, are also described in this section.

1. *The Present Invention*

15 The present invention is designed to allow voice and fax traffic to be carried over an IP network. This is more commonly referred to as IP Telephony. In addition to providing simple call termination, the present invention provides a host of enhanced features that enable IP Telephony capabilities to be used in a number of different applications. Since IP Telephony typically results in much
20 lower telecommunications costs, it can be used to reduce telecommunications costs between two branch offices. The reduced communication costs could instead be used to resell long distance calling to end users. A further application is enabling IP Telephony equipped terminals, such as multimedia PCs running group meeting software (such as but not limited to Microsoft NetMeeting), to
25 access traditional services, such as a call center. By combining IP Telephony with other technologies, such as the World Wide Web, integrated WWW call centers are possible. The present invention makes these applications possible by

acting as a gateway between traditional telecommunications networks and systems, and IP systems.

There are two main parts to each member of the present invention:

5 **Software Platform.** This is a suite of software components which provide IP Telephony services across an IP network. These components can run as executables or applications on an operating system (such as but not limited to processes running on Microsoft Windows). Each of the components of the present invention software is described in detail in the next section of this section.

10 **Hardware Platform.** In order to achieve connectivity with traditional telecommunications systems and networks, such as legacy phone systems (PBXs) and the public switched telephone network (PSTN), a specialized hardware platform, according to the present invention, is necessary. This hardware platform is used to run the present invention software components of the present invention that handle connectivity to traditional telecommunications devices, and
15 consists of a PC containing one or more specialized telephony hardware devices. In addition, standard PCs can be used to run other software components of the present invention that do not require a direct connection to the traditional telecommunications equipment.

20 Each member of the present invention consists of a hardware platform with specific traditional telecommunications equipment connection capabilities, as well as software configured for that connection. The present invention supports, but is not limited to, the following kinds of connections:

- Analog loop start (FXO, standard analog telephone lines)
- Analog station interface (FXS)
- 25 • T-1 robbed bit signalling
- ISDN PRI (T-1 and E-1)

In addition to the different types of connectivity supported, different members of the present invention provide different connection sizes, ranging from two line systems up to multiple T-1/E-1 systems. It is also possible to

license the present invention's software independently of the hardware platform, as well as to customize hardware platforms to achieve other capabilities of the VoIP system.

5 The main focus of this section is to explain the present invention in an example software implementation, and how it is used to deliver the desired IP Telephony services to the end user.

2. *Present Invention Components*

10 The present invention software is comprised of a number of independent software components, each of which provides some part of the overall functionality provided by the present invention. These components communicate with each other over the IP network that connects them. The different types of components are each discussed below. It is important to note that numerous instances of each type of component may be running simultaneously in different locations.

15 Each component of the present invention software is an independent application. Once installed, these applications can be executed in the same manner as any other application. Directions on interacting with the individual components are provided in the sections later in this document, with an overview of the interaction. Each of the components that comprise the present invention
20 software is described in the subsections below.

a. *Gateway Server*

25 Gateway Server 210 manages the specialized telephony hardware that is used for interconnection to the telecommunications equipment, and performs all bridging of calls between the traditional telecommunications network and the IP network to which the Gateway is connected.

The key responsibilities of Gateway Server 210 are shown in FIG. 6A and discussed below:

- Handling calls received on either the IP or traditional telephony networks (steps 605 and 610);

- Bridging calls between the two different kinds of networks (step 615). This involves voice compression, echo cancellation, control protocol transformation, and numerous other complex operations;

- Interacting with telephone users (step 620) using DTMF digits (and other methods of telephony signaling) in order to collect information such as the target number, user identification, password, and more;

- Interfacing with the routing system (step 625) (described below) in order to translate numbers and routing those numbers to the correct destination gateway, IP address, or telephone number (depending on the number dialed);

- Interfacing with the database server (step 630) in order to validate user accounts, apply charges, and keep records of all calls.

With multiple gateway servers, it is possible to do long distance PSTN switching over an IP network (step 635). This requires two gateway servers, one in the originating PSTN zone, and one in the terminating PSTN zone. This configuration is shown in FIG. 22A. The originating PSTN Zone 2210 is coupled to Gateway Server A 2220. Gateway Server A 2220 is connected to the Internet 2230. Gateway Server B 2240 is also connected to the Internet 2230. Gateway Server B 2240 is coupled to terminating PSTN Zone 2250. Calls are denoted by the dotted-lines which partially represent the flow of data through the system.

Of course, more than two gateways (i.e., gateway servers) could be connected via the IP network, resulting in a wider range of calling areas. Furthermore, it is not necessary to restrict calling to the local PSTN zone in which each gateway server is located. For example, a single gateway server in New York could handle all calls for North America. If the other gateway servers in the network were located outside of the United States and Canada, then this

would still represent significant cost savings over the direct dialed long distance rate from the international location.

b. Routing Server

5 Routing Server 230 is an optional component in the system that is used to route calls across the IP network, and possibly to a final destination on the telephone network. It is not a required component in order to be able to route calls. A gateway server can route a call based on local information maintained on the gateway server. However, the use of a routing server can facilitate the automatic exchange of routes between different gateway servers, and can in
10 general make the administration and configuration of call routing much easier.

The following are the key roles of the Routing Server 230, as shown in routine form in FIG. 6B:

- Allowing gateway servers to export their local routes for the use of other gateways (step 645);
- 15 – Transforming exported local routes into dynamic routes (step 650);
- Storing dynamic routes in a dynamic cache (step 655);
- Also storing disseminated routes and static global routes (step 660);
- Provide gateway servers with routes that are specifically intended for that gateway server (step 665). This allows all routes to be configured on the
20 server only, but automatically distributed to the appropriate gateways;
- Providing gateway servers with the ability to look up routes to a particular number (steps 670 and 675). For example, a gateway server could ask a routing server for routes to a number such as +1(703)555-1234. The routing server, if it has a route, would return it to the gateway server, which would then
25 use the route to complete the call (steps 680, 685, and 690). Hence, the routing server might indicate that there is a gateway server located at IP address 192.168.1.3 that can complete the call to +1(703)555-1234;

— Distribute E.164 translation data (as shown in FIG. 18 and discussed in detail below). An important process is how a string of digits entered by the user is translated into a standard, E.164-style address (e.g. 00117035551234 (+1(703)555-1234). The routing server facilitates the exchange of this data, eliminating the need to configure it on each gateway server.

FIG. 22B shows the example above in more detail. Routing Server 2265 is present in a network of Gateway Servers. Routing Server 2265 responds to a query from one of those Gateway Servers. A call originates from Telephone 2253 which is coupled to the originating PSTN Zone 2255, which is shown to be in the 212 area code. These specific digits and locations are purely illustrative and are not meant to limit the present invention. Originating PSTN Zone 2255 is coupled to originating Gateway Server A 2260 at internet protocol address 192.168.1.2. Gateway Server A 2260 is connected to the Internet 2230. Once a call is connected from the originating PSTN Zone 2255 to the Gateway Server A 2260, a request can be made to Routing Server 2265 over the Internet 2230. The request is for routes to a terminating Gateway Server B 2275 which is configured to handle telephone address +1(703)555-1234. Routing Server 2265 looks for matching routes according to its own features, described in detail below, and provides a route to the originating Gateway Server A 2260 that +1(703)555-1234 can be handled by a Gateway Server B 2275 at 192.168.1.3. Gateway Server A 2260 routes the call to Gateway Server B 2275. Terminating Gateway Server B 2275 is connected to the Internet 2230 and receives the routed call. Terminating Gateway Server B 2275 is coupled to terminating PSTN Zone 2280. Terminating Gateway Server B 2275 accessed the terminating PSTN Zone 2280 and places a telephony connection (i.e., a telephone call) to the requested telephone address. Terminating Telephone 2283 is connected to the terminating PSTN Zone 2280 and has been configured to answer calls at the address in question.

In a large deployment, and for reliability, it is possible to configure multiple routing Servers in the same network. The use of multiple routing servers

can provide both load balancing across multiple machines, and redundancy to protect against the failure of any particular routing server.

c. Database Server

5 The present invention software may, in many applications, need to manage the users of the system. When gateway servers are connected to the PSTN. PSTNs are accessible to/by the general public. Thus, in a VoIP system connected to the PSTN, it is generally necessary to restrict calling privileges to those users with valid user accounts. In addition, it is often desirable to bill for services, either post-paid or in real-time. Although the gateway servers could connect directly to a centralized database using Open Database Connectivity (ODBC), this approach would not be scalable. Conventional database packages support only a relatively small number of concurrent connections to the actual database. Furthermore, it would not be secure, particularly if the IP network connecting the gateway servers to the database is the Internet. In this environment, the actual database would need to be exposed to the Internet in order to provide access to gateway servers. It is often preferable to keep the actual databases protected behind a firewall.

10
15
20
25 Database Server 240 addresses these problems by providing middleware components that interface with numerous gateway servers, yet use a small number of connections to the underlying database. This allows very large numbers of gateway servers to have access to the data contained in a database accessible to Database Server 240. In one embodiment, Database Server 240 supports underlying databases in Microsoft Access format, as well as Microsoft SQL Server. The use of a standard underlying database package facilitates integration of external applications with the data that drives the present invention. The present invention is not so limited. Other types of databases and database formats may be used.

FIG. 22C shows one implementation of Database Server 240 in a network of gateway servers. A call originates through telephone 2284. Telephone 2284 is connected to originating PSTN zone 2285. Originating PSTN zone 2285 is coupled to Gateway Server A 2287. Gateway Server A 2287 received the call and communicates a login sequence with the user calling from telephone 2284. As shown in the figure, a user is prompted to enter an ID. The user's ID is passed on to Database Server 2289. Both Database Server 2289 and Gateway Server A 2287 are connected to the Internet 2230. Database Server 2289 looks up the user ID and determines if that user ID should be allowed access to the system, and if so, how much access. The present example simply shows that the user has some money in his account. Database Server 2289 sends an access message to Gateway Server A 2287. Gateway Server A 2287 either allows the call or provides a service notice and terminates the call. A routed call would proceed like in FIGs. 22A and 22B. A terminating Gateway Server B 2297 is connected to the Internet 2230. This gateway server routes the call to a telephone 2299 connected to terminating PSTN Zone 2297.

Even though only a small number of connections exist between Database Server 240 and the underlying database package, a License Agreement may require a separate license for each gateway server that either directly or indirectly uses the services of the database package.

More complex configurations are possible. As with the routing server, it is possible to use multiple database servers in a particular network. This is usually done to provide redundancy and protection against the failure of any particular database server.

d. Management System

Management System 250 is a GUI tool that is used to perform configuration and administration of the above present invention software components. It connects to the gateway, routing, and database servers remotely

via the IP network, and can be used to monitor the current state of those components, as well as to make changes to the configuration of those components. This architecture allows all of the gateway, routing, and database servers in the network to be configured from a single point. This is critical in environments where the actual servers may not be directly accessible. Furthermore, multiple instances of the management system can be run on different machines, allowing several administrators to jointly manage the network.

FIG. 162 shows a block diagram of the Management System Architecture 16200. Management System 250 includes a Management System Application Layer 16204, Management System Platform 16208 (software), and Management System Hardware 16212. Management System Application Layer 16204 includes Management System View 14604. This application layer provides the support for GUI Tools 255. In one example implementation where the software is written in an object-oriented program language such as C++, Management System View 16204 accesses lower-level modules (also called objects) through a Management System Document 14612.

Management System Platform 16208 includes Management Controls 14608, Management Client Objects 14616, Management Client 14620, Management System Document 14612, and ActiveX/OCXs 14624. These components provide the functionality of the management system by interacting with the other components of VoIP System 200. Management System Hardware 16212 includes the elements of Computer System 500 in the PC on which Management System 250 is executed.

FIGs. 160A and 160B illustrate a routine for Management System 250. The Management System Document 14612 is instantiated (step 16004) and connects to the Software Environment 14516 (step 16008). In step 16008, Management System 250 connects to the components of VoIP System 200 for which it is configured to connect, such as Gateway Server 210, Routing Server

230, and Database Server 240. The Management Client 14620 and Management Control Objects 14616 are instantiated (step 16010) and connect to each other (step 16014). An ActiveX/OCXs database 14624 containing commands, variables and functions is instantiated (step 16012). Management Client 14620 connects with Software Environment 14516 (step 16016), and thus to the components of Database Server 240 (step 16018). Management System View 14604 is instantiated (step 16020) and connects to Management System Document 14612 (step 16022). Management System View 14604 also connects to Management Client 14620 (step 16024).

Management Controls 14608 are instantiated (step 16026) and connect to Management Client Objects 14616 (step 16028)

From its connection with Software Environment 14516, Management Client 14620 determines if an update to ActiveX/OCXs 14624 is required (step 16030). Any update to ActiveX/OCXs 14624 is implemented in both Management Client Objects 14616 and Management Controls 14608 (steps 16032 and 16034).

Management System 250 is discussed in further detail below in the Key Concepts section.

e. Provisioning System

Provisioning System 260 provides GUI Tools 265 that allow access to the data managed by Database Server 240. Management System 250 provides access to Database Server 240, and allows Database Server 240 to be configured, but it is the role of Provisioning System 260 to actually allow interaction with the data contained in the databases. The following are the main tasks of Provisioning System 260:

- Adding, removing, or modifying the accounts of system users;
- Configuring and maintaining group records for users, and account records for users; and

– Reviewing call detail records (CDRs) created by the gateway servers and stored within the components of Database Server(s) 240.

Provisioning System 260 can be operated remotely, just like Management System 250, and can be used to administer multiple Database Servers 240. In some applications, the provisioning system may not be required, since external applications can directly modify the data contained in the underlying database. The provisioning system is provided as a tool for those users who don't require the flexibility of a custom external application.

FIG. 161 displays a block diagram of Provisioning System 260. Provisioning System 260 includes a Provisioning System Application Layer 16104, Provisioning System Platform 16108 (software), and Provisioning System Hardware 16112. Provisioning System Application Layer 16104 includes Provisioning System View 14504. This application layer provides the support for GUI Tools 265. In one example implementation where the software is written in an object-oriented program language such as C++, Provisioning System View 16104 accesses lower-level modules (also called objects) through a Provisioning System Document 14508.

Provisioning System Platform 16108 includes Provisioning System Document 14508. Information is obtained through Software Environment 14516 from Database Client Subsystem 14512. These components provide the functionality of the provisioning system by interacting with the other components of VoIP System 200. Provisioning System Hardware 16112 includes the elements of Computer System 500 in the PC on which Provisioning System 260 is executed.

FIG. 159 illustrates a routine for Provisioning System 260. Provisioning System Document 14508 is instantiated (step 15904) and connected to Software Environment 14516 (step 15906). Provisioning System Document 14508 connects to Database Server 240 through Software Environment 14516 (step 15908). Provisioning System View 14504 is instantiated (step 15910) and

connects to Provisioning System Document 14508 (step 15912). Users are provided information from Database Client Subsystem 14512 via GUI Tools 265 which utilize Provisioning System View 14504 (step 15914).

f. Utility Components

5 The present invention includes a number of utility components that are useful for the initial configuration as well as diagnosis of present invention software components. These components do not necessarily play a large role in the day to day operation of the system. These components are as follows:

10 *Installation Program.* This program will install the present invention software components that one can choose on to a particular machine. One normally need only run this once, and only if one purchased the software independently of the hardware platform.

15 *Express Configuration.* The Express Configuration (or ExpressConfig) program is designed to allow one to set up a basic configuration for present invention software components in a short period of time. The most common values for advanced features are automatically configured, other features are automatically configured based on a simplified set of information that one provide during Express Configuration. After the process is complete, the configuration of the system can be fine tuned using the management system.

20 *Control Panel.* The control panel is an application that allows one to monitor, start, and stop the gateway, routing, and/or database server(s) running on the same machine as the Control Panel utility. It is useful for diagnostic purposes, but it has no capability to change any configuration parameters.

3. *Other Implementation Details*

a. *Components and Nodes*

As explained in the previous section, there are a number of different types of components in the present invention software. In addition, it was also explained that there may be many instances of the same type of component, running on different machines. Using the gateway server example, there might be one gateway server in New York, and another gateway server in Los Angeles. These two gateway servers might switch calls over an IP network connecting them, allowing inexpensive long distance between New York and Los Angeles.

The term "node" will often be used to describe an instance of a particular type of component. For example, both of the gateway servers in the previous example are distinct nodes, even though they are both components of the same type.

It is also possible to run instances of different components on the same machine. In this case, each different component is in effect a different node. There is actually no difference in operation regardless of whether components run on the same machine, or on different machines, except for the utility components which only operate on the local machine. Therefore, when nodes are discussed in this document, it will be assumed that they could always be running on the same machine. However, it is not possible in some cases to run two or more instances of the same component on the same machine. As an example, a gateway server and a database server could run on the same machine. However, in some cases two routing servers could not run on the same machine.

b. *Operational vs. Administrative Components*

The present invention software consists of three types of components.

Operational Components. This refers to the gateway server, routing server, and database server. These software components provide IP telephony services to the end user, and are not interacted with directly by the administrator

of the system. Operational components are always running, and run in the background with no user interface. Operational components are similar to other software services; they can be started or stopped at any time, and only provide their services while started. Operational components do not, however, run as operating system services – they are software applications that run in the background on an operating system. Operational components can be administered at any time, even while stopped, as long as they are running.

Administrative Components. This refers to the management system(s), and provisioning system(s). These components are applications that can be used to interact with the operational components. The management system interacts with all three operational components in order to perform configuration changes; the provisioning system interacts only with the database server in order to manipulate data in the database.

Utility Components. This category of components was defined previously, and provides programs that operate directly on the local machine.

c. *Configuration*

All present invention software components manage configuration in the same way. The configuration is stored in the operating system's registry, under a key. Each individual product component will store configuration in subkeys of this key. When a present invention software component first starts, it will read its configuration from the registry. After startup has been completed, components will not in general re-read the registry. If one makes configuration changes to the present invention software directly by editing the registry, one will in general need to restart the component for which one is editing the configuration in order for the configuration change to take effect.

In general, one will not need to interact directly with the configuration data stored in the registry. The main purpose for providing this information is to

allow one to locate the configuration data so that one can ensure that it is backed up with the rest of an installation.

d. Identification and Authentication

5 The security model employed in the present invention employs two main ideas. The first is the idea of identification and authentication, and the second is that of access control. Identification and authentication, described in this section, describes how one component of the present invention, usually a server of some sort, determines the identity of another component that is requesting some service or data. Once the requesting component has been identified and authenticated, access control is then used to determine if the component with that identity should indeed be granted access to the requested data or service.

10 Identification and authentication are performed through the concept of a security domain. This concept is central in trying to configure the security subsystem of any present invention component. Stated simply, a security domain is just a collection of members that are part of that domain. Each member has a unique name within that domain, and has some authenticating information that can be used by other members in the domain to verify that it is indeed a valid and legitimate member of the domain. The domain will also have a unique name. When displaying the full name of a member, the DOMAIN:MEMBER syntax will be used, where DOMAIN is the name of the security domain, and MEMBER is the name of the member within that domain.

15 An present invention security domain should not be confused with a operating system's domain. The two types of domains are not related in any way. Making two present invention components part of the same security domain does not make them part of the same operating system domain.

25 Each instance of an present invention software component, previously defined as a "node", must be a member of at least one security domain. This is always required, since all components of the present invention either connect to

some node to administer and manage them (which must of course pass identification and security checks), or else they are connected to by administration/management components (in which case the connecting component must be identified).

5 It is possible for a component to be a member of multiple domains. However, there are not many situations in which this capability would be necessary.

Also, recall that two nodes running on the same machine (for example, a gateway server and a database server) are not considered to be the same node. Therefore, despite running on the same machine, the two components could be given independent, different identities.

10 Lastly, it should be noted that two distinct nodes can be given the same identification. Typically, this is used to simplify administration by assigning all components running on a given machine a common identity. However, it is generally recommended that different components of the present invention be given identities even if they run on the same machine.

e. Access Control

As described in the previous section, the security model in the present invention is comprised both of identification and security, and access control. Identification and authentication are a prerequisite to access control. The role of access control is to determine what level of access should be provided to a remote node. For example, just because a gateway server can determine the identity of a management system, it does not imply that the gateway server should automatically allow that management system to administer it.

20 Access control is required in numerous places in the system. The main places where access controls is required are as follows:

The gateway, routing, and database servers need to determine whether a particular management system should be allowed to perform administrative operations.

5 The database server needs to determine whether a particular provisioning system should be allowed to modify the data in the database.

The routing server needs to determine whether a gateway server should be allowed to retrieve call routing data.

The database server needs to determine whether a gateway server should be able to query the user database, or insert CDRs.

10 Although there are access control features between gateway servers, they are handled using a different mechanism to allow additional features (such as billing of the remote gateway) which are not possible using the access control system applied in other areas of the present invention.

15 Access control in the system is accomplished through an access control list associated with the particular service or resource that can be accessed. This access control list specifies the identifiers (such as DOMAIN1 MEMBERA) of the members that are permitted to access the service or resource with which the access control list is associated. Each entry in the access control list contains not only the identities of the members that should be associated with that entry, but
20 also information concerning the level of access that the identified members should have. For example, the routing server allows control over whether or not a gateway server can export routes for the use of other gateways.

f. Management

25 The management system component is used to perform the management and administration of the gateway, routing, and database servers, as described previously. This administration and management is fully remote, and can be performed from anywhere in the network providing that there is an IP network connection to the node being administered.

Having explained the security model employed by the present invention components, it is now possible to explain the basic mechanism by which the management system functions. Each operational component has a management server – it allows management operations to be performed by a remote management client. The client in this case is the present invention management system and is discussed previously. In order to manage a remote system, the management system first establishes a connection to the operational component. After this connection is established, the management system provides identification and authentication information. The operational component checks the information provided by the management system, and uses this information to determine whether or not the connecting management system should be granted to access to administer it. If the connecting management system cannot be identified and authenticated, or it does not have a matching entry in the Management access control list on the target operational component, the connection will be terminated.

Once the connection is successful, the management system can be used to administer the operational component. How to use the management system to perform this administration is described in detail below.

g. Licensing

In order to use an present invention software component, a software license for the component being used must be available. This license can be provided upon the purchase of the software.

According to the present invention, licensing is performed via license keys. A license key is an encoded string that contains information regarding a particular license. Each license key provides a fixed number of licenses to a particular component of the present invention software. License keys are only valid for use by the customer to which they were issued. Furthermore, license keys are dependent on a hardware key, also supplied at the time of purchase. This

prevents one's license key from being illegally used by others that may somehow obtain access to this license key. The hardware key attaches to the parallel port of one's computer, and allows the license key to be interpreted correctly.

5 In one example, license keys are required only for the gateway server component of the software. One gateway server license must be acquired for each line in the system. (However, a single license key may contain 10,100, or more licenses, so a separate license key is not required for each line). Since multiple license keys may be used simultaneously, it is possible to expand a system (providing that adequate hardware is installed) simply by purchasing additional license keys. This provides the flexibility of acquiring a small number
10 of licenses to begin with, and scaling up the number of licenses as traffic requires.

There are two main modes of licensing that can be used. These licensing modes are as follows:

Standalone Mode. In standalone mode, each gateway server is required
15 to have its own hardware dongle. A license key is required for each gateway server, and must contain sufficient licenses for the use of that gateway server. This model is simple to implement, and allows complete independence between Gateways.

Server Mode. In server mode, a centralized license server manages
20 licensing for all components in the network. When a component first starts up, it will contact the license server to receive the necessary licenses. With this configuration, all license keys are configured on a single machine, and only one hardware key is required.

Time-limited, demonstration license keys are also be available to facilitate
25 evaluation of the licensing.

3. *Installation and Setup*

This section describes an example installation and setup of the present invention. This section is divided into three main sections:

5 *Hardware Installation and Setup.* This section details how to install the hardware used by the gateway server to interface with the telephone network, and how to configure the drivers for that hardware.

Software Installation. This section describes how to install the various present invention software components.

10 *Express Configuration.* This section describes the Express Configuration program, and how it can be used to configure a default, initial configuration for the present invention.

15 In one example, the present invention is a turnkey systems. The hardware, driver software, and software of the present invention installed and pre-configured for a particular environment. Alternatively, software can be sold separately. The information in this section pertains only if the present invention software was acquired separately, or if the software needs to be re-installed.

1. *Hardware Installation and Setup*

20 This section discusses the installation and setup of hardware devices used by the gateway server to interface with the telephone network. Telephony hardware devices are not required in order to operate the routing server, database server, or any of the other components of the present invention software.

 There are two main types of hardware devices that are used by the gateway server. These hardware devices are:

25 Telephony interface hardware devices. These hardware devices are used to connect the PC on which a gateway server operates to a telephone network. A wide variety of such boards are supported, in order to enable connectivity to a variety of hardware devices. Dialogic Corporation manufactures the telephony

boards used by the gateway server. The following models of Dialogic hardware are supported:

- D/21H and D/41ESC in Gatenet 3 mode (2 line analog loop start)
- D/41ESC in normal mode (4 line analog loop start)
- D/160SC-LS (16 line analog loop start)
- D/240SC-T1 (24 channel T-1 robbed bit, 23 channel ISDN PRI)
- D/300SC-E1 (30 channel ISDN PRI)
- MSI/SC-R (8, 16, 24 channel analog station interface)

It is possible to use multiple boards to achieve higher numbers of supported lines within a single chassis. It is also possible, in some cases, to mix boards of different types.

DSP devices. These devices perform digital signal processing (DSP) on telephone signals using high performance, embedded DSP chips. These devices are required in configurations where a large number of lines need to be handled by the system. Boards from two hardware vendors are supported: Analogic Corporation, and Brooktrout Technology. The following boards are supported:

- Analogic TAP-801 (for backwards compatibility only)
- Analogic TAP-802 (22/33-port DSP board)
- Analogic TAP-804 (24/36-port DSP board)
- Brooktrout TR-2001 (24-60 port DSP board)

a. Installing Dialogic Telephony Boards

Before installing the Dialogic Telephony boards, make sure that switches on the board are set correctly. Switches vary from board to board, follow the instructions (if any) corresponding to the board being installed listed below.

- On a D/41ESC, set/push both switches on SW2 to the bottom level.
- On a D/160SC-LS, ensure the board is "on-hook" (see SW4 on Quick Install Card).

- On a D/240SC-T1 board, set "loopback" off - push switch down (see SW500 on Quick Install Card).

All Dialogic telephony boards that are to be installed in the system should be inserted into the machine as per the Quick Install Guide shipped with the boards. The boards must be firmly placed in to the slot. Because the Dialogic boards (with the exception of the D/21H) are full length boards, it is important to ensure that there are no obstructions when inserting the board, or else the board might become loose at a later time.

Once the boards have been inserted, one must next confirm that each Dialogic board has a unique ID. The ID for each board is set by way of a rotary switch located on the top of the board, near the front of the machine into which the board is inserted.

b. Installing DSP Boards

If one has DSP boards, they should be installed into the PC in the same manner as the Dialogic Telephony boards, and following the directions in the users guide supplied with the DSP board. One will also need to following certain board-specific directions when installing DSP boards:

- For the Analogic TAP-801/802, check the rotary switches at the bottom of the board near where the board plugs into the ISA bus. There are two small rotary switches that control the I/O address of the board. By default, this I/O address is 0x240, which is specified using a '2' on the first dial, and a '4' on the second dial. If one is installing multiple boards, he or she must ensure that each board has a unique I/O address. 0x250 usually works for the second I/O address.
- For the Analogic TAP-804, no jumper/switch settings are necessary.
- For the Brooktrout TR-2001, no jumper/switch settings are necessary.

c. Installing the SCbus cable

If one is installing DSP boards, he or she will also need to install an SCbus cable in order for the system to operate correctly. If one is not using SCbus (i.e. using Gatenet 3), then the SCbus cable is not required. One will need an SCbus cable that has at least one drop for each board that is installed in the system. There are several important notes regarding SCbus cables below:

- Even if one has a Dialogic board in the system that is not planned to be used, one must still connect the SCbus cable. Failing to do so may prevent the Dialogic service from initializing.
- Some Analogic TAP-802 boards have a notched SCbus connector that is backwards. In no event should one twist the SCbus cable to fit into this notch! If a SCbus cable is in fact notched (meaning that one is unable to normally insert it into the connector on the Analogic TAP-802), then a user must either acquire a non-notched SCbus cable, or remove the notch.
- The SCbus cable should run flat across the connectors of all ISA boards in the system. If one is using a PCI DSP board, then one will require a longer SCbus cable. However, the cable should still be flat—there should be no twists in the cable.

Symptoms of an incorrectly connected SCbus cable include not being able to initialize the Dialogic boards even though they are correctly detected, and not hearing any audio even though the gateway server answers successfully.

d. Installing the Dialogic Drivers

After the hardware boards have been installed and the PC has been restarted, it is necessary to install the Dialogic DNA3 drivers in order to verify that the boards are correctly installed and to enable the gateway server to make use of the boards. The following procedure is used to install the Dialogic drivers.

The gateway server operates preferably with Dialogic DNA3, but is not limited to Dialogic DNA3. Previous versions of the Dialogic drivers (DNA,

DNA2, 98.02, 97.08, or others) can be used. The Streams environment is not required in order to install the DNA3 drivers.

Step 1: Obtaining the Dialogic drivers

First, it is necessary to obtain the Dialogic DNA3 drivers. There are several sources from which the Dialogic DNA3 drivers may be obtained:

A CD may be purchased provided with a purchase of the present invention. The Dialogic driver diskette may have accompanied boards, if one purchases the hardware separately.

Array Telecom Corp's FTP site at <ftp://ftp.arraytel.com>. The Dialogic DNA3 drivers can be found in the directory /Version3/Drivers/DNA3, and is supplied under the filename DNA3.zip.

Dialogic Corporation's WWW support site at <http://support.dialogic.com>. One will need to follow the links on the web page, and provide some information in order to be able to download the DNA3 drivers from the Dialogic web site.

If one downloaded the Dialogic software either from Dialogic, one will need to uncompress it into a directory from which one can install it.

After one obtains the Dialogic drivers, one will need to run the installation program. If running the installation program manually, run SETUP.EXE from the directory containing the Dialogic drivers. If installing from an original Dialogic DNA3 CD, the setup program may be started automatically when the Dialogic DNA3 CD is inserted into a CD-ROM drive.

A welcome screen 2302 appears once the setup program is executed, as shown in FIG. 23. Welcome screen 2302 contains text concerning general installation parameters as well as Next button 2304 and Cancel button 2306. The implementation of these button bullets is well known in the relevant art. A person skilled in the relevant art will recognize these button bullets and their functionality. Button bullets appear throughout the figures which follow and have not always been labeled, however their function and implementation conform to those button bullets already labeled and described above.

A user information screen 2402, as shown in FIG. 24, can appear if the Next button 2304 was pressed in the welcome screen 2302. User information window 2402 contains user name field 2410 and company field 2412 as well as Next button 2404, Cancel button 2406 and Back button 2408.

5 FIG. 25 shows setup options screen 2502. The setup options screen 2502 allows one to select which components should be installed. Several options 2510 can be available. Here, four are shown: typical, complete, compact, or custom. For instance, one selects "typical" unless there are special requirements, or unless one wishes to select ISDN, in which case custom can be selected.

10 If one selected custom and selected ISDN, one will be prompted to see if ISDN should be installed. If one selected custom, one will also need to indicate whether the ISDN Package should be installed. One should answer yes to this question if using the Dialogic D/240SC-T1 or Dialogic D/300SC-E1 in ISDN mode to connect to telephony equipment, or intends to make use of ISDN in the near future, and does not wish to reinstall the drivers later. (Installing ISDN drivers if ISDN is not in use will not cause problems).

15 FIG. 26 shows the protocol selection screen 2602. If one chooses to install ISDN, he or she will be prompted to specify from the list in the Components panel 2610 which ISDN protocols should be installed. One should select the protocol that corresponds to the ISDN switch to which the system will be connecting.

20 If installing ISDN, and connecting to a device that is normally connected to a DMS 100, AT&T 4ESS, or AT&T 5ESS switch, select the NT1 protocol. This protocol emulates the network side of an ISDN connection.

25 FIG. 27 shows the location selection screen 2702. One is prompted to enter the location to which the Dialogic drivers will be installed. It is recommended that one use the default options for the installation location. A user can click on Next to install to a location other than the default.

FIG. 28 shows the file selection screen 2802 which displays the selected install options for final review before installation. A user will then be provided with an opportunity to confirm that the selections made match particular requirements. Once the information provided in the display is checked, click Next button to begin installing the actual files.

Once the installation completes successfully, a user will be given the option of viewing the release notes, and/or running the Dialogic configuration manager. Regardless of whether or not one chooses to read the release notes (not necessary in most cases), one should choose to run the Dialogic configuration manager.

It is recommended that one not reboot your system immediately after the installation, since this is not necessary, and one will need to reboot after installing other drivers and software.

Next, a dialogic configuration manager screen 2902, shown in FIG. 29, automatically attempts to detect which boards are installed in a PC. In general, boards will be correctly detected providing that they are installed correctly. Screen 2902 can show detected boards.

Once boards have been directed, a user may need to make configuration changes depending on your system setup. The changes necessary for operation with the gateway server, based on a user's board type, are as follows. A user may also need other changes depending on the exact environment in which the software is being configured.

FIG. 30 shows Protocol Configuration screen 3002 For Dialogic D/240SC-T1s and Dialogic D/300SC-E1s operating in ISDN PRI mode, one will need to configure the ISDN protocol used on the device. Select the device, click on properties, go to the interface tab, and select the ISDN protocol from one of the protocols shown in the list. One will need to do this for each ISDN device installed in his or her system.

If running in Gatenet 3 mode with the Dialogic D/41ESC or D/21H, one will need to make changes to the Dialogic configuration. These changes are explained in the next section. Once configuration of the Dialogic drivers is complete, the Dialogic service can be started.

Step 4: Obtaining and Installing Dialogic Gatenet 3

It is only necessary to install Gatenet 3 for a low density software only configuration. For hardware DSP processing boards the Gatenet3 must be installed for the software to work correctly.

Dialogic Gatenet 3 is essentially a set of firmware from Dialogic that runs onboard the D/21H and D/41ESC boards. The purpose of the Gatenet 3 package is to provide hardware echo cancellation, allowing real-time, full duplex conversations to take place in an IP Telephony environment.

Installing Gatenet 3 overwrites some files from the Dialogic DNA 3 installation. As with the Dialogic DNA 3 software, it is necessary to obtain the Dialogic Gatenet 3 drivers first. These drivers can be obtained from the following locations:

A compact disk having the DNA 3 software on it;

The CD;

A Dialogic CD containing the Gatenet 3 drivers.

Array Telecom Corp's FTP site at <ftp://ftp.arraytel.com>. The Dialogic DNA3 drivers can be found in the directory /Version3/Drivers/GN3, and is supplied under the filename GN3.zip.

Once the Dialogic Gatenet 3 drivers are obtained, one runs SETUP.EXE in order to install them. One is then greeted with Welcome screen 3102 shown in FIG. 31. Welcome screen 3102 contains text with general installation information and Next and Cancel buttons.

When a user selects "Next", he or she may receive a warning screen, indicating that the Dialogic Gatenet 3 package overwrites certain files and registry entries, and cannot be uninstalled. This means that if a user needs to

uninstall Gatenet 3, he or she will be forced to reinstall the Dialogic DNA 3 drivers to have a stable configuration.

A user will next be requested to enter user and company information. Enter this information, and click on Next.

5 The installation options will be confirmed. Verify that the installation options are correct, and click next. One should ensure that the directory to which the Dialogic Gatenet 3 drivers are installed matches the directory to which one originally installed the Dialogic DNA 3 drivers. Select "Next" once the installation options are verified. Setup will then proceed to copy the necessary files.

10 After optionally reviewing the release notes following a successful installation, a user will be given the option of rebooting his or her machine. It is recommended that a machine be rebooted at this time, prior to starting the Dialogic service.

15 After a machine has rebooted, one runs the Dialogic Configuration Manager in order to complete the configuration of the Gatenet 3 package. Run the Dialogic Configuration Manager by going to the Start Menu \ Programs \ Dialogic System Software \ DCM.

20 One can make the following changes using the DCM. The following changes must be made under the properties of each D/21H and/or D/41ESC installed in the system:

Under the *Misc*, set Gatenet=ON

Under the *Dialogic Bus* tab, set BusType=NONE.

25 In addition, one should edit *gatenet.prm*. This file is found by default in C:\Program Files\Dialogic\Data. A user may use any standard text editor, such as Notepad or Write, to make these changes. However, the resulting file should be saved in a text format.

One will find a block in *gatenet.prm* that reads:

parm[246] = <value>

parm[247] = <value>

In the above, <value> will actually be a number. Change the value so that the section reads:

parm[246] = 240

5

parm[247] = 240

Once the above changes are made, save the file. One will need to restart the Dialogic service in order for the changes to take effect.

Step 5: Starting the Dialogic service

10

When the Dialogic drivers are installed, they are not activated until a user selects to actually start the service. Starting the service also involves verification by the Dialogic drivers that it is possible to start the Dialogic boards.

15

In order to start the Dialogic drivers, use Start Menu \ Programs \ Dialogic System Software \ DCM to start the Dialogic Configuration Manager. If a red button with "STOP" written in the middle appears in the toolbar, then the Dialogic service is already running. Otherwise, a Green button should be present in the toolbar, which a user can press to start the Dialogic service. Depending on a particular board configuration, the Dialogic service can take from between 15 to 60 seconds in order to start.

20

Once a user has confirmed that the Dialogic service starts correctly, he or she may wish to configure it to start automatically whenever the machine starts. In order to do this, open the Service menu in the DCM, and select Startup mode. One can change the Startup mode between manual and automatic in this menu.

e. *Installing Analogic Drivers*

25

This section will describe an example installation of the Analogic drivers. These drivers are required to support the TAP-801/802/804 DSP boards manufactured by Analogic Corporation.

The Analogic drivers should be installed if the hardware configuration includes Analogic DSP boards. Gatenet3 users can choose not to install Analogic drivers

Step 1: Obtaining the Analogic drivers

5 The first step in the installation is obtaining the Analogic drivers. One should obtain the correct version of the Analogic drivers. For example, one can select the v1.3 drivers.

10 There are sources from which one can obtain the Analogic drivers: a compact disk having v1.3 drivers, or Array Telecom Corp's FTP site at <ftp://ftp.arraytel.com>. The Analogic v1.3 drivers can be found in the directory \Version3 \ Drivers \ Analogic1.3, and is supplied under the filenames TAP-H.EXE and TAP-D.EXE.

Step 2: Uncompressing the Analogic drivers

15 After one has successfully obtained the v1.3 Analogic drivers, they may be decompressed so that they can be installed. A CD can contain uncompressed versions of TAP drivers in addition to the compressed version, so one does not have to follow these directions if installing from CD.

20 One should uncompress both the TAP-H and TAP-D drivers separately. In order to uncompress the TAP-H drivers, run TAP-H.EXE from the location where it was downloaded. One will need to provide the name of a directory to which the TAP-H drivers will be decompressed. It is recommended that a user select C:\TEMP\TAP-H, so that he or she can easily remove the temporary files afterwards. Once a user clicks on OK, the files should be uncompressed, and ready to use.

25 FIG. 32 shows display screen 3202 for uncompressing to a temporary directory. Display screen 3202 contains a directory selection field for determining the location to which files will be uncompressed and buttons to uncompress (or unzip), run wizip, close, about and help. The same procedure should be followed for TAP-D, except that one must uncompress to a different

directory. C:\TEMP\TAP-D is recommended. As with TAP-H, this is a temporary directory that can be removed after the installation is complete

If installing from a CD, the uncompressed versions of the TAP-H and TAP-D drivers can be found, respectively, on the CD. A user can use this in place of C:\TEMP\TAP-H and C:\TEMP\TAP-D in Steps 3 and 4 below.

Step 3: Installing the TAP-H driver

After the drivers have been uncompressed, one can install the TAP-H drivers by running SETUP.EXE from the directory into which he or she uncompressed the TAP-H driver. The TAP-H driver installs the host libraries necessary for an application (such as an example of the present invention software) to interact with the DSP board. One should install this driver first, before TAP-D.

After one runs SETUP.EXE, he or she will be greeted with welcome screen 3302 as shown in FIG. 33. Welcome screen 3302 contains general installation instructions and next, back and cancel buttons.

FIG. 34 shows location selection screen 3402. One is allowed to enter the location to which the TAP-H drivers are installed by selecting the "Browse..." button. One should install to the default location, since some of the TAP utilities are highly dependent on pathname.

After a user has selected the directory to which the TAP drivers are installed, he or she will be asked for the program folder that should be created to hold the TAP-H installed programs. A user may accept the default, or modify it if a different folder name is preferred.

Following the selection of the folder to which the TAP-H drivers are installed, the copying of the necessary files will commence. After the copying is complete, a user will be advised of the completion, and asked to log off and log back on in order for the changes to take effect. Preferably, one should not log off since he or she will need to reboot the machine following the installation of the TAP-D drivers.

The Analogic TAP-H drivers modify environment settings for the currently logged on user. If the gateway server is to be run using a different account, the settings made by the TAP-H driver will need to be moved from 'User Variables for <current user>' to 'Systems Variables' under My Computer \ Properties \ Environment. The variables that should be manually copied between the sections are 'TAP_800' and 'TAP_DIAG_DIR'.

Step 4: Installing the TAP-D driver

The TAP-D driver is an driver used in addition to TAP-H in order to use an Analogic TAP-801/802/804 DSP board. The TAP-D driver installs the software to initialize with and communicate with the DSP board. The utilities installed by TAP-H are not usually used until TAP-D has been installed.

The installation for TAP-D is very similar to that for TAP-H. SETUP.EXE must be run from the directory to which the TAP-D driver was uncompressed. It will prompt a user with welcome screen 3502 as shown in FIG. 35. Welcome screen 3502 contains general installation information and next, back and cancel buttons.

As with TAP-H, a user will be allowed to select the directory into which the software will be installed. A user should select the same base path as you did for TAP-H, and click Next. As with TAP-H, a user will also be allowed to select the folder into which the driver will be installed.

When the installation is completed, one will be advised to run *drv_install* and then to reboot the system. One should not follow this procedure precisely, since it does not allow one to verify that the board was in fact installed. In order to confirm correct installation of the driver, one should follow the procedure below.

Open a new Command Prompt, using Start \ Programs \ Command Prompt.

Change to the drive where the TAP-D driver (e.g. C:) is installed.

Change to the directory where the TAP-D driver (e.g. C:\opt\analogic\tap800) is installed.

Change to the subdirectory containing the board driver. (e.g. CD host\drv).

Run install.bat by typing 'install'.

5 This batch file should install the driver, and configure it. One should see the following message:

C:\opt\analogic\tap800\host\drv>drv_install

tap installation

Querying the system for available resources.

10 This may take a while...

Probe found board 0, a Tap802, at port 0x240

If one sees the line above, or a line similar to the above, it indicates that the board has been located successfully.

15 After one confirms that the drivers have been correctly installed, it is recommended that one reboot his or her system.

f. Installing the Brooktrout Drivers

This section will describe the installation of the Brooktrout drivers. These drivers support the TR-2001 series of DSP boards manufactured by Brooktrout Technology.

20 The Brooktrout drivers should be installed if the hardware configuration includes Brooktrout TR-2001 DSP Boards.

Step 1: Obtaining the Brooktrout drivers

25 The first step in the installation is obtaining the Brooktrout drivers. There are sources from which one can obtain the Brooktrout drivers: a compact disk having the Brooktrout drivers.

Array Telecom Corp's FTP site at <ftp://ftp.arraytel.com>. The Brooktrout drivers can be found in the directory /Version3/Drivers/Brooktrout, and is supplied under the filename TR2001.zip.

Step 2: Installing the Brooktrout drivers

Run the file TR2001SDK31.EXE that was obtained in Step 1 above. This will start the installation procedure. The Brooktrout installation will be automatically uncompressed to a temporary location. After this process is complete, the setup program will be started automatically.

FIG. 36 shows the license agreement review and acceptance screen 3602. The Brooktrout TR-2001 SDK setup screen will appear. One can click Next at this screen to proceed. One should then review the license agreement.

One will be presented with the option as to where the software should be installed. A user should accept the default location in most cases. After one selects the location to which the software is to be installed, the copying process will commence. A series of MS-DOS/command prompt windows may appear during this process. This is normal.

After the installation is complete, a user will be requested to reboot and should reboot his or her computer at this time.

2. *Software Installation*

Once the necessary hardware and associated device drivers have been installed, a user will be able to proceed with the installation of the present invention software. When installing the software, one will be able to install as many instances of the software as one wants. In a fuller example, one must have the appropriate numbers of license to actually be able to use the software.

In one example of the present invention, software components are all installed from the same installation program. One may select the specific components that are to be installed when the installation program is run. This section will step through the installation of the present invention software, and explain the decisions made in the process.

Before a user begins, the user obtains the present invention software distribution. Once a user has obtained and uncompressed installation distribution,

he or she is ready to begin installation. Prior to beginning the actual installation, a user should be sure to have the following information available: the Customer ID for an organization, the Product ID, and the License Keys corresponding to the software that is licensed.

5 When a user has the information listed above, he or she is ready to begin the installation process.

 From a directory containing the uncompressed installation distribution, run a program SETUP.EXE. This will commence the installation process. One will be greeted with welcome screen 3702 shown in FIG. 37. Welcome screen 3702 contains general installation information and next and cancel buttons.

10 Following welcome screen 3702, one is greeted by information screen 3802 that requires one to enter information about him- or herself, his or her company, and information on a license sheet. Information screen 3802 is shown in FIG. 38.

15 After the information screen, one will be given a choice of setup types. Two options will be presented, complete, and custom. FIG. 39 displays the installation option screen 3902 where these options are displayed in a panel. Complete will install all components of the present invention software on the local machine. Custom will allow the user to select which components should be locally installed. Generally, unless one is setting up a test environment, it is best to select the components of the software that one will actually require or use.

20 FIG. 40 shows selection screen 4002. Selection screen 4002 displays various installation options and buttons for back, next and cancel. If installing the gateway server, one should additionally select the type of gateway server configuration that should be installed. There are three choices, corresponding to the different hardware configurations of a gateway server. They are explained below.

25 Analogic TAP-80x. Select this option if one is using Analogic DSP boards with a system.

Gatenet 3: Select this option if one is using only Dialogic D/21H or D/41ESC boards, and has the Gatenet 3 package installed.

Brooktrout TR-2001: Select this option if one is using any of the family of Brooktrout TR-2001 boards.

5 After selecting a DSP configuration, one will be asked for a program folder to install the software to. Once the software has been installed, one will be able to access it in this folder. Normally, the default folder is acceptable. Once a user clicks Next, the copy process will begin. FIG. 41 displays the installation of a startup menu group in setup and installation screen 4102.

10 Once the copy process is completed, a user will be given the option of viewing a readme file that accompanies the software. This file will usually contain the most up to date information about the software.

15 It is not necessary to reboot a machine after installing the present invention software. Once the software has been successfully installed, a user is ready to move on to the configuration of the present invention software.

3. *Express Configuration*

This section describes an Express Configuration (ExpressConfig) program. ExpressConfig can be used to generate an initial configuration for present invention software. It should be used when the software is first installed, 20 or when the software needs to be completely reset, since it will replace all configuration data. Any configuration changes made prior to running ExpressConfig may be lost.

25 ExpressConfig performs configuration of the gateway server, routing server, and database server components of the present invention. It is capable of performing any or all of these components when it is run, so it is possible to reset and configure the database server without affecting the configuration of the gateway server or of the routing server. Nonetheless, it is necessary to exercise

care when running ExpressConfig to ensure that configuration is not unintentionally overwritten.

a. Running Express Configuration

FIG. 42 shows the express configuration welcome screen 4202. Express Configuration is installed when a user installs the present invention software. Clicking Start \ Programs \ present invention \ ExpressConfig will start the ExpressConfig program. Once the ExpressConfig program has been started successfully, a user will be presented with a welcome screen, shown below. One will be able to navigate backwards and forwards in the configuration process by using the tree view in the left pane of the window. Alternately, a user can step through each configuration step in a normal sequence by simply using the Next and Back buttons.

The ExpressConfig program is divided into three main sections corresponding to different components of the present invention that can be configured. There are some aspects of configuration that are common to all three components that are set up using ExpressConfig, these common aspects are discussed in the next section.

With all components, there are three possible options with respect to how ExpressConfig will operate on that component. The three choices are presented to the user in the gateway server screen 4302 of the ExpressConfig screen shown in FIG. 43.

An inset panel 4304 displays the available installation options. Not all of the choices may be shown if one is installing for the first time. An explanation of the choices follows:

– Do not configure the <component name> on this machine. One should select this machine if one has existing configuration for the given component that you do not want to overwrite, or if one does not wish to set up the component at this time. This is the default option if configuration already exists.

– Perform configuration of the <component name> on this machine. This option will allow one to configure the given component on the local machine. With this option, items not specified during ExpressConfig will not be overwritten; in other words, only the areas of configuration handled by ExpressConfig will be affected. This is the default option if configuration does not already exist.

– Reset the configuration of the <component name> and perform new configuration. This option is similar to performing configuration normally with the exception that any existing configuration will be automatically overwritten with default values. This option should be used carefully, but may be useful if one has completely corrupted your existing configuration.

Once a user decides which of the above is appropriate, he or she can proceed to entering the actual configuration information.

Note that ExpressConfig is a local utility. One cannot use it to remotely configure components located on different machines. This can be accomplished using the management system after the present invention has been installed and set up using ExpressConfig.

b. Common Setup Information

This section describes the setup information that is common across the gateway server, routing server, and database server. The setup information described in this setup pertains to common things, such as identification, security, and management. As described above, ExpressConfig will step through a sequence of information screens. The following are the common information screens that a user can interact with for all three components.

(1) Identification

An identification panel is used to configure identification and authentication information that will be used by a particular component to verify the entities that connect to it and use its services. Although in principle a user could configure a complete list of all entities on each component, in practice this is not necessary. The following list shows, for each component, what kinds of entities will connect to it and use its services:

- Gateway Server: Management system for access to management data.

Note that other gateway servers do not use the same identification/authentication mechanism.

- Routing Server: Management system for access to management data.
- Gateway Servers for access to routing data.
- Database Server: Management system for access to Management data
- Provisioning System for manipulation of database.
- Gateway Servers for user management and CDR collection.

With the above list in mind, it is only necessary to configure identification/authentication information for the entities that will connect to the component that you are configuring. For example, when configuring identification/authentication information for a gateway server, a user need only enter identification/authentication information for any management systems that may connect.

When setting up identification/authentication information using ExpressConfig, it is very important that a user at a minimum provide a configuration that allows for a management system to connect. Otherwise, one will may be able subsequently to connect using the management system in order to modify the configuration of the system, and may be forced to re-run ExpressConfig. FIG. 44 shows the identification screen 4402 in ExpressConfig.

The Security Domain Settings panel 4404 contains a security domain name field, a server identity field, a listing window of current security domain members, and a sub-panel for adding additional security domain members.

Specifically, the first two parameters a user needs to supply are *Security Domain Name* and *Server Identity*. In essence, one is naming the security domain that the server will be a member of, and is providing the name/identity of the server within that domain. The name of the security domain will need to match the name of the security domain of the entities that the server being set up will interact with. For example, if one were setting up a gateway server that was managed by a management system in domain TESTDOMAIN and which connected to a routing server in domain TESTDOMAIN, then one must use the same domain name – TESTDOMAIN. The identity that one gives the server within the domain should be unique. Any components that the server being configured connects to will need to have identification/authentication for the identity that one assigns the server.

Once a user has decided on the security domain name, and the identity of the server being set up within that domain, one will need to provide identification/authentication information for the other members of the domain. As mentioned above, one only needs to provide this information for entities that make use of the services provided by the server being set up. In order to add to the list of identities that can be identified and authenticated, enter the member name and IP address in the boxes provided in a panel or other GUI input. A new entry will appear in the list in the middle of the panel. To remove one of the entries, simply select it and click Remove. The principles behind this process are explained in more detail in the next section, when the identification/authentication model is elaborated upon. Advanced features (identity wildcarding, the use of netmasks) are also explained in the next section.

(2) Management Access

FIG. 45 shows the management access settings screen 4502. The Management Access settings screen contains a gateway server management access settings panel 4504. Panel 4504 contains a list of clients which can access the gateway server and a sub-panel for adding additional clients. Thus panel 4504 allows one to specify the identities of the nodes that should be granted management access to the server that is being configured. Any node listed on management access control list will be able to connect to the server being configured, and make changes to the configuration of the server. Therefore, it is important that the management access control list be used carefully in conjunction with the identification/authentication settings to ensure that unauthorized access is prevented.

The Management access control list is reasonably simple, because ability to manage is a binary ability – all nodes identified on the list are capable of managing, and any node not identified will not be allowed to perform any management functions. The means of adding and removing from the management access control list is similar to the identification/authentication list. To add a new entry, enter the member name of the node that is to be added, and click on the Add button. Similarly, to remove an entry, select the entry and click on Remove. It is assumed during ExpressConfig that any member a user adds will be in the security domain configured during the identification/authentication step of the ExpressConfig process.

(3) License Information

In one example, components of the present invention require licenses in order to provide services. The use of license keys allows one to ensure that he or she is always in compliance with his or her license agreement, and that use of the software does not exceed the licenses that have been purchased. During

ExpressConfig, a user can add license information, and an initial license key. A user may add additional license keys as he or she purchases them.

In one example, although licensing information can be configured on all of the gateway server, routing server, and database server, only the gateway server actually makes use of licensing information. However, licensing capabilities are included on all of the servers to ensure flexibility in creating licensing arrangements that meet customer requirements.

FIG. 46 shows the license information screen 4602. There are two main pieces of information that are used in setting up licensing using ExpressConfig. The first is customer/product identification information, shown in panel 4604. One will need to provide a customer ID and product ID. Both of these pieces of information can be provided on a license agreement.

After a user has entered the basic identification information, he or she enters licensing information into licensing information panel 4606. There are two modes in which licensing can function. The first is standalone mode, where keys are locally configured on each server. If operating in this mode, then enter the license key supplied with the system. If, on the other hand, one is operating in client/server mode, he or she provides the IP address of a primary license server (and a secondary license server, if any). More information on setting up licensing is contained in the next section.

c. Gateway Server Setup

This section describes the setup that is specific to the gateway server component of the software. There are three panels used to configure gateway server specific configuration data.

(1) Hardware Setup

The hardware setup panel provides a user with the means to quickly setup the hardware configuration of the system. It automatically creates configuration entries for the hardware devices were selected, and creates matching lines and line groups (explained in more detail in Section 5) based on selections.

FIG. 47 shows hardware setup screen 4702. Hardware setup screen 4702 contains two panels: DSP configuration panel 4704 and hardware devices installed panel 4706. There are two main questions that are to be answered in order to perform the hardware configuration of a gateway server. Firstly, the type of DSP configuration that is present in a system is identified. Secondly, a user indicates the number of Dialogic boards installed in the system that are to be used by present invention.

If the Dialogic boards are being used in the same chassis as a gateway server with an application other than the gateway server, the board number that is opened following ExpressConfig completion ExpressConfig uses boards starting from 1 and progressing upwards.

The DSP configuration of panel 4704 is selected by picking one of the three available DSP options. One must select the DSP configuration that corresponds to the version of the software that is installed and to the hardware configuration of an actual system.

For the installed hardware device list of panel 4706, there are a number of hardware devices listed. In front of each device, there is an editable field showing how many devices of that type are assumed to be present. One needs to fill in the numbers of each device that are present in the system. Note that one may not be able to select Gatenet 3 mode D/41ESCs and D/21Hs if it was indicated that a hardware DSP board based configuration is used. Similarly, a user can only select quantities of the above boards if a user selected a Gatenet 3 DSP configuration.

Additionally, a user must indicate the type of installed MSI/SC boards, if any, as well as whether the D/240SC-T1 boards are configured in robbed bit mode or ISDN PRI mode. A user's selection matches the actual hardware configuration of the machine, and the configuration of the Dialogic drivers performed in previous steps. If a user does not install any MSI/SC boards, or D/240SC-T1 boards, then the respective settings can be ignored.

(2) Call Routing

A call routing panel is used to allow one to configure the list of routing servers that the gateway server being configured will access in order to obtain routes. There are numerous settings that can be used to customize how a gateway server interacts with a routing server; the default values are used for any servers set up during ExpressConfig. Furthermore, significant local routing information can be set up on the local Gateway; however, this information, if any, should be entered using the management system. A Call Routing page shows a standard list. FIG. 48 shows call routing screen 4802. Call Routing screen 4802 contains a routing server list panel 4804. In order to add a new routing server to the list, enter a name for the routing server, provide its hostname or IP address, and click on the Add button. To remove a configured routing server, select it in the list, and click the Remove button.

(3) Database Access

The database access screen 4902 is shown in FIG. 49A. The database access screen 4902 contains a Database Server List Panel 4904. The Database Server List Panel 4904 allows one to set up the list of database servers to which the gateway server being configured will connect. This page is identical to the Call Routing page described above, except that it is database servers being set up. In addition, ExpressConfig will create default settings for the database services (user management, CDR collection) based on the servers that you configure.

d. Routing Server

The routing server has a single unique panel in its configuration, which is the client setup panel 4950 used to configure the list of clients that will be allowed to connect to and use the services of the routing server. The client setup panel 4950 is shown in FIG. 49B. As mentioned previously, these clients are gateway servers, however, other systems designed to obtain similar information can be connected through this configuration panel.

The setup of the routing servers client access control list is identical to setting up the management access control list. Gateway servers listed in the client access control list will be able to connect to the routing server, export routes, query for available routes, and receive disseminated routes. The exact permissions of a particular client can be modified on a client-by-client basis using the management system.

e. Database Server

There are two panels of information that need to be set up using ExpressConfig for the database server. These panels are Client Setup and Server Login Setup and are described below.

(1) Client Setup

Client setup for the database server is identical to the setup of the client access list for the routing server, described above. Just as with the routing server, the database server client access list can be manipulated at a later time using the management system.

(2) Server Login Setup

FIG. 50 shows Server Login Setup screen 5002. The Server Login Setup screen contains fields 5004 for to provide information that is used to interface with the underlying database containing user records, call records, and so forth. The information provided on this screen is used by the Database Server 240 to establish and maintain the connection to the underlying SQL database 245.

There are four key pieces of information that are used. The first is the Database Source Name (DSN). One creates an ODBC data source (using Start \ Control Panel \ ODBC) with the corresponding name prior to actually running the database server. One follows the directions below in the section on Database Table Setup to ensure that the database is setup with the appropriate tables. Although a user does not need to do this prior to running ExpressConfig, it should be done prior to running the database server.

When supplying a DSN in order to allow the database server to connect to the database, it may also be necessary to specify login credentials (a username and password) that are required to access the database. If a username and password are required, then a user should enter them in the supplied fields. Finally, a user should supply the name of a log file that will be used to record problems in interacting with the underlying database.

4. Database Table Setup

If a user will be using the database server component of the present invention software, it is necessary to ensure that the underlying database is correctly configured prior to starting the database server. There are three ways that a database can be configured for operation with the database server:

The database tables can be created and configured manually, using the schemas provided in Appendix A. This is not preferred unless a user intends to modify an existing database to fit the schema provided in Appendix A.

For example, if one is running Microsoft Access, a user can use an empty database file that contains all the required tables, with no initial data in them. The database file can be installed as part of the installation of the database server or located on a CD in uncompressed form. The present invention is not limited to this database format.

In another example, if a user is running Microsoft SQL Server, he or she can use a utility installed with the database server. This utility will connect to the specified data source, and create the necessary tables. One may use this utility by running setupSQLDB.exe from the directory to which the database server was installed, or by running an execution program from a CD. The present invention is not limited to this database server.

5. Additional Setup Issues

This section describes additional setup instructions that should be followed to ensure maximal performance of the present invention in one example implementation.

a. Screen Savers

PCs contain a number of screen savers that can be enabled after a certain period of inactivity. However, it should be noted that some of these screen savers can utilize 100% of the available CPU resources, since they perform complex graphical rendering. It is recommended that the screen saver be disabled on machines running the present invention software, or that a simple, non-3D screensaver be selected.

The present invention software runs at elevated priority to handle excessive CPU use by lower priority processes. However, it is still recommended that care be taken not to overload the machine running present invention software with other applications unless the performance of the present invention software on that machine is not critical (e.g. management system).

b. Find Fast / Virus Scanners

Microsoft Office by default installs a file indexer, called Find Fast. This allows you to quickly locate Office documents. However, it periodically indexes the entire hard drive, which can degrade the performance of the entire system while it is running. See the above note regarding CPU utilization under screen savers.

The same is true for Virus Scanners, which may be configured to periodically read all files on the hard disk, scanning them for viruses.

4. Getting Up and Running

This section explains how to get present invention system up and running, once the basic installation and setup described in the previous section has been completed. In this section, one will learn how to start the system, and use the management and provisioning systems to interface with the system. One will also learn how to perform operations that are common across different components of the present invention software, such as configuring security and management access control.

This section is divided into the following sections:

Starting the Software – Describes how to run operational components.

Control Panel – Shows how a Control Panel can be used to monitor and interact with the operational components.

Using the Management System – Explains the basic principles necessary to use the management system to interact with a running operational component.

Service Control and Logging – Shows how to use the management system to control an operational components service state, and how to view and log events generated by that operational component.

Security Domain Configuration – Explains how to manipulate the security domain settings for an operational component.

Management Access – Describes how management access to a particular operational component can be restricted, so as to prevent unauthorized access and modification.

5 *Licensing* – Covers how licensing works, and how license keys can be entered and used with present invention software components.

Using the Provisioning System – Explains the use of the provisioning system to interact with data managed by the database server.

Troubleshooting Service Startup – In the event that a service cannot be started, explains possible causes.

10 1. **Starting the Software**

15 The first step in using the present invention software is to start it. In a Microsoft Windows operating system example, once the present invention software is correctly installed, a program group can be found under the Windows Start menu containing all present invention software components. From this entry in a start menu, one can choose which components to start. FIG. 51 is a screenshot of screen 5102 and contains an exemplary start menu/programs menu 5104. Start menu/programs menu 5104 contains links 5106 to the software of the present invention.

20 In order to get started, one will generally need to run at least two components; one operational component, such as the gateway server, plus the management system, which is used to connect to the operational component in order to monitor and administer it.

25 Once started, the operational component will display a small icon in the system tray. Through this icon, one can perform very simple operations, such as starting and stopping the service provided by the operational component that is running. The small letter in the upper left corner the icon indicates which operational component is being interacted with; G represent the gateway server, R the routing server, and D the database server. Right clicking on the icon will

bring up a small popup menu. From this popup menu, one will be able to see the about dialog, which displays program version information. One will also be presented with the ability to start or stop the service, unless the service is already starting or stopping. One can use the Close item to terminate the service completely. Once terminated, a service cannot be interacted with in any way until it is started up again. FIG. 52 is a screenshot 5202 of the service menu popup menu 5204.

Based on the ExpressConfig process that one went through, one should be able to load all of the operational components that one configured, and use the icon in the system tray to start those services. The state of a service is indicated by the color of its icon. Red represents a stopped service; yellow represents a service that is either starting or stopping, and green represents a service that is started. If one chooses to start a service, and the service fails to start, then the color of the icon will return to red. This indicates a configuration error that one will need to correct using the management system prior to being able to start the service. The last section in this section explains how one can troubleshoot such a condition.

2. The Control Panel

When the Control Panel is started, a Control Panel screen 5302 is provided as shown in FIG. 53. The Control Panel screen provides a means of interacting with locally running operational components, without having to go through the Management System 250. It is provided for convenience, and has only limited capabilities for interacting with the operational components. It is capable of interacting with the Gateway Server(s) 210, 220, Routing Server 230, and Database Server 240. The Control Panel is started the same way as other present invention components, using Start \ Programs \ present invention \ Control Panel.

Tabs 5304 along the top of the screen allow one to select which component's panel one wishes to interface with. For example, the gateway server tab 5306 is shown. This tab contains sub-panels for starting and stopping gateway servers, configuring service levels, and showing event logs.

5 Once the desired tab 5304 has been selected, there are several things that one can do:

Run – This button can be used to run the service, if it is not already running.

10 Connect – This button will connect to the service, assuming that it is running. If the service is not running, then one will not be able to connect successfully. Interacting with a component requires that one first connect to that component.

Start/Stop – Once one is connected, one will have the option to start or stop the service (unless it is starting/stopping already).

15 In addition to these controls, one will also be provided with diagnostic messages generated by the operational component that one has selected, once one has connected to it. These diagnostic messages are used to indicate information that can be useful in determining whether or not the system is functioning correctly. Underneath the event log window, which displays these diagnostic messages, one will also have several controls that allow one to control the display of these messages. Pressing the 'Clear' button will remove all messages currently displayed in the event log window. Pressing 'Lock Display' will prevent any new messages from being added, until one press 'Unlock Display' (which appears only when the display is locked). This allows one to read a particular set of events in more detail, without worrying that they will be repositioned or disappear because of new events that occur. New events will be collected while the display is locked, but one will be able to scroll up and down through the existing events. Finally, a message level selector allows one to choose what level

20

25

of events one are interested in. The lower the level selected, the more messages will be displayed

At the bottom of the interface, service-specific information will be displayed. The information displayed is service dependent, and is as follows:

5 Gateway Server – A list of the current sessions is displayed, showing the lines that have active calls, and the services that are active on those lines.

Routing Server – A list of the connected clients (gateway servers) is displayed, along with their IP addresses and connect times. This window is useful for determining whether or not a gateway server is actually able to connect.

10 Database Server – A list of the connected clients (either gateway servers or provisioning systems) is displayed. The IP address and time of connection is displayed. As with the routing server, this is useful for determining if gateway servers are actually able to connect to the database server.

3. *Using the Management System*

15 This section introduces the management system, used to perform all administrative and management aspects of running an present invention network. The concept of how management is performed was discussed above; in this section, how to apply those concepts in order to actually manage operational components is discussed.

20 The management system is a Microsoft Windows GUI application used to connect to and administer operational components of the system. In a Windows example, it can be started once the present invention software is installed using Start \ Programs \ present invention \ Management system.

a. Initial Configuration

Once one has started the management system, one will need to do two things in order to be able to use it to manage your network. Each of these steps is described below.

5 *Security Domain Configuration.* Security domains are configured during ExpressConfig for each operational component that one configured. In order for a management system to connect to an operational component, it must share a common security domain with that operational component. In addition, from a security domain standpoint, it is also necessary that the operational component
10 be able to authenticate the management system as a valid member of that domain.

In order to configure the identity of the management system within a security domain, one should select the "Security" item from a "Local Identification" menu. This will bring up the Security Domain Configuration screen 5402 as shown in FIG. 54. From this screen, one can list the security domains that the management system is part of.

15 Clicking on the Add button contained in the Security Domain Configuration screen 5402 will allow a new security domain to be added. When adding a security domain, the only relevant information is the Domain Name, and the Member Name within that domain. Other fields are not required, although
20 one can provide a description for future reference.

In order to change a domain, it is necessary to first remove the existing domain, and a new domain with the same name. The only reason to do this would be to change the member name of the management system within that domain. This is not frequently required. Selecting a particular security domain
25 will provide additional information about that security domain.

It is possible to configure multiple security domains on a given management system. However, each security domain must have a unique name. It is not generally necessary to do this, unless a particular management system needs to configure operational components located in different security domains.

b. Remote System List

A management system is capable of administering numerous remote systems. In order to facilitate this process without needing to enter addressing information for a system each time it is necessary to manage that system, the management system maintains a remote system list that is permanently stored. The remote system list appears (by default) on the left side of the management system window, and shows all remote systems that have been entered.

There are several ways to manipulate the remote system list. The first is through the "Remote Systems" menu. When this menu is selected, the following additional items appear:

Add New System: Upon selecting this option, one will be provided with a dialog panel 5502 as shown in FIG. 55. This provides a Server Settings panel 5504 for creating a new system in the remote system list. The dialog panel 5504 needs to be used to add each operational component that will be managed from the management system.

One needs to provide the necessary information in this dialog that pertains to the system being added. The fields required in one example are as follows:

System Name: An identifying name for the system that is being added. This name is for local identification purposes only, and does not have any impact on connecting to and/or managing the remote system being configured. It does not need to match, and is in no way related to, a member name within a security domain. However, it may be easiest to use a nodes security domain member name in this field.

System Type: Indicates the type of node, gateway server, routing server, or database server.

IP Address: IP address of the node that you wish to connect to.

Port Number: UDP/IP port number that the server is running on. You should normally accept the default setting, which is filled in automatically when you select the correct system type.

Description: This field allows you to enter arbitrary text that will help to identify this system. It is not required, but can be used to record information for reference purposes, e.g. "Located in New York, USA".

Once a user clicks on OK, the new remote system will be added, and will appear in the remote system list.

– Remove. This command removes the currently selected remote system from the managed system list. One may only select this command if an entry is currently selected in the remote system list.

– Connect/Disconnect. This command will establish or terminate a connection to an operational component. One must have a remote system selected in order to perform this command. Once connected, one will be able to administer the operational component.

– Properties. This option allows one to configure the settings for an existing remote system, and requires that the remote system is selected in the remote system list.

The second way of performing these commands is by using the mouse, by right clicking in the remote systems list. If one right clicks on a remote system, one will be able to connect/disconnect from it, view its properties, or remove it. If one right clicks an empty area in the remote systems list, one will be given the option to add a new remote system.

c. Connecting and Administering

Once one has configured at least one entry in the remote systems list, corresponding to the operational component(s) that one wishes to administer, one will need to connect to the operational component. One can do so by issuing the connect command, as described above, or simply by double clicking on the remote system to which one wishes to connect.

Once connected, an additional tree of options will appear beneath the remote system that a user connected to, in the remote system list. This tree will

allow a user to navigate between the options that are available to configure. The exact makeup of this tree varies from one type of component to the next, and also depends on the configuration of a particular operational component.

If one has difficulty connecting to the remote system, there are several possible causes. Regardless of the cause, a message box will appear indicating that a connection to the remote system could not be established.

If this dialog appears immediately, then the remote system that you are trying to administer was contacted, but it rejected your attempt to connect to it. This is due to one of two reasons:

A user could not be identified and authenticated by the remote system. A user should check the security domain name, and member name of the management system, in order to ensure that they are correct. Then, a user should confirm that the remote system that he or she is trying to administer is in the same security domain, and that the security domain on that remote system is capable of authenticating the user. This process is explained in the section below, Security Domain Configuration. As long as a user entered the member name and IP address of the management system when using ExpressConfig, this should be taken care of.

The management system was identified and authenticated, however, it had no entry in the management client access list. Configuring the management client access list of an operational node is discussed in the sections below, however, as long as the member name of the management system was entered during ExpressConfig, it will have an entry in the management access list.

If this dialog takes several seconds to appear, then the remote system could most likely not be contacted. This may be due to an error in the IP address of the remote system, a temporary network failure, or a firewall that is blocking traffic. It may also be caused if the operational component that a user is trying to connect to is not running, or if a user selected the wrong type of system when setting up the remote system entry in the remote system list.

The Remote System List window 5602 in FIG. 56 is in expanding tree view form and shows the remote system list after a connection has been established, and the connected node has been expanded. The tree view is navigated in the same way as a tree view presented by many file system viewing applications. Nodes in the tree can be expanded and collapsed by clicking in the boxes beside them. Nodes with a '+' in the box beside them can be expanded. Nodes with a '-' beside them can be collapsed, or else do not have any nodes beneath them. Double clicking on the label for a node will bring up a window corresponding to that node (the contents depend on which node was selected). Changes to configuration are made using the windows that appear. Future sections describe in detail how to interact with individual components using the management system.

4. Service Control and Logging

When one uses the management system to connect to an operational component, one screen will automatically appear. This screen is the Service Manager screen 5702 as shown in FIG. 57. The Service Manager screen 5702 is present on all operational components. It contains a Service Status panel 5704, an Event Logging panel 5706, and a Message field and list window.

Several options are available from this screen. Each of these options is discussed below.

a. Service State Control

First, the service state can be modified by either starting or stopping the service, assuming that the service is not already in the state of starting or stopping. The current service state is also indicated. Clicking on Start/Stop will cause the service to be started or stopped, respectively.

Event logging can also be configured from the Service Manager. Clicking on "Event Log" will bring up the event logging configuration dialog 5802 as

shown in FIG. 58. Event logging allows events to be logged to disk for subsequent review. Event logging incorporates a circular logging feature that allows the amount of disk space consumed by the log to be managed. An event logging configuration panel 5804 allows one to enable or disable the logging of events to disk. Normally, logging events to disk is unnecessary unless one is trying to diagnose a problem. If one enables event logging, he or she will have the following options:

Debug Level. Selects the level of events that will be logged to disk. Any events that are of a lower level than the level selected will not be stored in the log. This can help to only log events of the relevant level to disk, so as to maximize the amount of useful information in the log file.

Maximum File Length. Log files are circular; once a particular log file fills up, the next log file in sequence will be truncated to zero length, and used to store log information until it too is full. The maximum file length parameter is used to control how full a log file is allowed to get before the next log file in the sequence will be used.

Log File List. This list provides the names of the log files that will be used to record events. These log files will be used in a circular pattern, using the process described above. One can use the Add/Remove buttons to add or remove log files from the list, or can set the current log file (i.e. the log file into which the next event will be placed) by clicking on the Set Current button. Normally, the default log filenames will be sufficient.

Clicking OK will make changes to logging configuration permanent; pressing Cancel will abort any changes that one made.

Finally, the Service Manager window provides a means to monitor events that occur on a particular operational component. Any events that occur on the operational component that is currently connected to will be displayed in the window. However, events may be filtered by priority level. The default debug level can be used to control what level of events will be displayed. It is not

normally necessary to view debug events, and since each message consumes bandwidth, it is generally sufficient to leave the debug level on the default setting of only displaying warning-level events or higher. The Advanced button allows one to configure event display levels on a type-by-type basis, and is only required for advanced diagnostics. Finally, two buttons are provided in order to clear and lock/unlock the display.

5. *Security Domain Configuration*

One of the concepts explained herein is that of security domains. The concept of a security domain is central to understanding how identification and authentication operate in an present invention network.

Nodes in an present invention network need to be in a common security domain in order to exchange services. In order for two nodes in the same security domain to be able to authenticate each other, they need to have authentication information that allows them to authenticate each other. There are many possible forms that could be used for the exchange of authentication, such as a password. The present invention software is constructed so as to support a variety of methods for identification and authentication. However, at this time, only a single model for authentication, based on member lists, is supported.

Security domain names, and member names within a security domain, are case sensitive. It is imperative that attention be paid to case when entering security domain names and member names.

From the management system, a top level node in the tree of the gateway server, routing server, and database server is the "Security and Identification" node. Double clicking this node will bring up the Security and Identification screen 5902. FIG. 59 shows this setup screen. From this page, one can interact with the Security Domains panel 5904 and configure the operational component that is being managed. The Security Domains panel 5904 contains an expanding tree view window and buttons for adding a domain, adding a member, removing

either a domain or a member, and a properties button for displaying the properties of a selected domain or member. A detailed description of these buttons follows.

a. Memberlist Domains

In a memberlist domain, authentication checking is performed by way of a static list of members, and corresponding IP addressing information. When a node providing some service receives a new connection requesting access to that service, it will check the identity of the requesting node, and determine if the IP address of the requesting node matches the corresponding entry in the database. If there is a match, then the identity of the requesting node will be accepted. Otherwise, the identity will be rejected, and the connection will progress as if the requesting node had no identity. (Normally, this would result in a denial of access to the resource).

In FIG. 59, it can be seen that a security domain can be expanded, by clicking on the '+' next to the name of the domain. Expanding the domain causes the members of the domain to be shown. The "Add Member" and "Remove" buttons can be used to add and remove members from the member list. Clicking on "Properties" allows the settings for a member of the domain to be configured.

When a new member is added, or an existing member is edited, Properties screen 6002 will appear. FIG. 60 shows the Properties screen 6002 that is displayed when a new member is added or an existing member is edited. This screen contains two panels 6004.

In FIG. 60, the Identification panel 6006 is shown. Three settings can be configured for a member of a security domain. The first is the name of the member. When a request is received from a remote node, it will present its name. The name that is configured in this dialog must match the name presented by the remote system when it connects. The second piece of information required is the privilege class to be assigned to the remote system. Privilege classes are not currently used by existing components in the software, so using the default value

of zero suffices. The third field is a description of the node. This is purely for informational and reference purposes, and does not affect the functioning of the software in any way.

Once these basic settings have been completed, Addresses panel 6106 can be displayed. It allows the IP address authentication information to be configured. The Addresses panel 6106 is shown in FIG. 61.

In Addresses panel 6106, a user can configure a list of IP addresses from which the named member of the domain will be allowed to connect. If the IP address from which the member connects is not within one of the ranges specified in the list, then the identity of the connecting member will be rejected. The list of addresses is manipulated in the usual fashion. Each entry in the list has three pieces of information:

- (1) The IP address from which the member is allowed to connect;
- (2) The netmask controlling which IP addresses can be connected from.

For example, if the IP address field is 192.168.1.33 and the netmask is 255.255.255.0, then connections will be accepted from 192.168.1.0 through 192.168.1.255;

- (3) The port number from which the member is allowed to connect. The default value of zero allows a connection from any port, and should be used unless it is necessary to override this setting.

Note that multiple address ranges may be configured. It is only necessary for a connecting member to match one of these address ranges in order to be accepted. Thus, if the same name is assigned for convenience to two management system nodes, one IP address entry could be entered for each node, allowing them both to identify themselves as the authorized management system.

It is possible to enter a '*' (asterisk) for the name of a domain member. If such a wildcard entry exists, and no other entries in the list matches the member name presented by a connecting node, then if the addressing information for the wildcard entry matches that of the connecting node, the identity of the

remote node will be accepted as valid. For example, configuring an entry with a name "*", and an IP address 192.168.1.0 netmask 255.255.255.0 would allow all nodes from 192.168.1.0 through 192.168.1.255 to connect, and their identities to be accepted no matter what is presented.

6. *Management Access*

In addition to configuring the identification and authentication information on a particular operational component, to have administration by a remote node, access should be granted to the services provided by that operational component. This section describes how to configure access to the Management 'server' built into each operational component.

All operational components have a root level node in their configuration trees titled: Management Access. Double clicking on this node will bring up a Management Access dialog. This dialog presents the list of nodes that currently have management access. For each entry in the list, the name of the entry, description of the entry, identification corresponding to the entry, and current status of the entry are shown. The status simply shows whether or not management is current enabled for that entry. Entries in the list are called "clients" – in effect, they are "clients" of the internal management "server" of each operational component. These clients are the management systems that are authorized to connect. FIG. 62 shows Management Access screen 6202.

When a new entry is added to the Management Access List within Management Access screen 6202, or an existing entry is modified, the Client Configuration screen 6302 shown in FIG. 63 is displayed. This dialog allows properties to be set up for an entry in the management access list. In the dialog of FIG. 63, it is possible to configure the General panel 6304 and Identities panel 6306. In the General panel 6304, the first is the name of the access entry. This name is for informational and logging purposes, but must be unique. The second is a description of this entry, which is also for informational purposes only. A

flag allows an entry to be temporarily disabled. If this flag is set, then management systems with an identification matched by this access entry will not be allowed to manipulate the management system.

In the Identities panel 6306, a list of identities that is accepted for this management access entry can be configured. The purpose of this list is to specify those nodes that are able to access the management system using this access entry. This list enables multiple distinct remote identities to share a single client configuration. The purpose of this is mainly for future expansion. For example, if a management access entry was extended to support the configuration of exactly which screens could be accessed, it would be tedious to duplicate this configuration for each management system that needed to be configured. Therefore, allowing multiple identities allows the access entry to be shared amongst multiple management systems.

When new identities are added, or existing identities are modified, a security identification dialog pops up, allowing a user to enter the identity being added or modified. The Security Identification screen 6402 is shown in FIG. 64.

In the Security Identification panel 6404 of the screen 6402, the type allows a user to select whether the identity being configured is a remote node or a security group. Security groups are not presently accessible directly by the user, so foreign entity should generally be selected. In addition, a user should specify the domain name and member name which comprise the identity of the connecting node. The domain name may be selected from one of the security domains that is configured on the operational component that the management system is connected to, and the member name should match the name of a member that can be authenticated using the authentication information discussed in the previous section. In addition, it is possible to wildcard the member name by entering a single asterisk (*), which allows any valid member of the domain to be recognized under this access entry. Lastly, a minimum privilege level can be configured, in which case, the requesting node must have this level of

privilege, or higher, in order to be granted access. Privilege levels are configured for each domain member when setting up a security domain.

7. *Licensing*

The software licensing model used by present invention components has been described hereinabove. This present section provides further details on how to setup software license keys, and elaborates on standalone and client/server modes of operation so that it is possible to select the correct choice for a given configuration. Each of the two licensing models, and how to configure the software for operation under either of those two models, is given below. In one example, only the gateway server requires license keys to be entered. However, the process for managing licenses is common to all present invention software components that require licensing, and in alternative embodiments of the invention, can apply in different licensing models.

Regardless of the mode of operation selected for licenses, an operational component that uses licensing can have a top-level management tree node named "Licensing." Double-clicking on this node will bring up a License Manager screen 6502 as shown in FIG. 65. This screen contains a License Manager Setting panel 6504 and a List of encrypted license keys panel 6506.

There are several common pieces of information that are used regardless of the licensing mode actually in use. The ExpressConfig program normally fills in these fields, but it may be possible to change these fields at some point. Any changes to the configuration of the License Manager can only be performed by clicking Change Settings button in the License Manager Setting panel 6504, and making changes using the dialog that appears. The common settings are as follows:

Mode. This allows the License Manager to be changed in terms of its mode of operation. A user should select the value for this parameter first, since certain parameters cannot be edited unless a user is in a particular licensing mode.

Customer ID. A user must enter the customer ID assigned to his or her organization in this field. Each license key issued is keyed to a customer ID. This information is normally printed on a software license certificate, issued when a user purchased the software.

5 *Product ID.* The product ID is a serial number which identifies the product. Each license key issued is keyed to a Product ID, and will only work with this product ID. Like the customer ID, the product ID is printed on a software license certificate.

10 The remaining settings are determined by the licensing mode that a user selects. Each of the available licensing modes is discussed in the following subsections.

a. Standalone Mode

15 In standalone mode, all licensing on a particular node in the system is performed independently of all other nodes. All license information required by that node is held locally. In standalone mode, each node should have its own hardware dongle in order to decrypt license keys.

20 License keys can be added and removed from the main License Manager window. This option is provided once standalone has been selected as the mode of operation. A user should select standalone operation, and then accept the changes prior to being allowed to enter license keys. Two simple buttons, Add and Remove, allow a user to add and remove license keys. When a user clicks “Add” and enters a license key, the license key will be verified, and it will either be accepted or denied. If it is accepted, it will appear in a License Manager window, and information associated with that license key will be display. If, 25 however, it is denied, it will not be recorded. If a license key is repeatedly denied, then a user should check the following:

Ensure that the license key is typed in exactly as issued. There are no brackets, apostrophes, or quotation marks in the license keys themselves, so if a

license key is provided within such delimiters, a user should not type them in when entering the license key. For example, if an E-mail reads:

Your license key is "ABCDEFG1234567"

then it should be entered as ABCDEFG1234567, without the quotation marks.

Case is significant when dealing with license keys. Therefore, ABCDEFG1234567 should not be typed in as abcdefg1234567.

Ensure that the customer ID that a user has entered is correct, and matches the customer ID on the license certificate or otherwise issued to the user (the customer ID may not always be printed on the certificate depending on the restrictions on the key, but can generally be found there).

Ensure that the product ID that a user has entered is correct, and matches the product ID on the license certificate or otherwise issued to the user. (As above, the product ID may not always be printed on the certificate).

If a hardware dongle is used (hardware dongles are generally used; time-limited, demonstration keys may not use a hardware dongle), one must ensure that the dongle is correctly connected to the computer, and that it is connected to the correct kind of port on the computer.

If using a limited license restricted to a particular IP address, ensure that the IP address of the machine that a user is installing on matches the IP address used to issue the license key.

License keys are time limited, and may have expired. Ensure that the system time and date are accurately set, and that the license key being entered is presently valid.

b. Client/Server Mode

In client mode, a centralized license server determines the licenses available to a particular operational component (a client of the license server). This mode of licensing is suitable in environments where a number of operational components are owned and operated by a single organization. It reduces the necessary administration since licenses can be acquired for multiple operational components and centrally configured, rather than needing to enter license keys on each individual node. However, license keys issued for use by a license server can only be used by a license server, and thus, any operational node that no longer operates with the license server will be unlicensed, and unable to provide services. For example, an organization deploying numerous gateway servers in different areas would benefit from the use of a license server. However, if those gateway servers were independently owned and operated, the gateway server owners may not want to use a license server. This is since their licenses would effectively be pooled, and would not be separable should a gateway server owner decide to leave the association.

Operation in client/server mode requires the presence of at least one centralized license server. In addition, it is possible to configure a backup server, so that if the primary server becomes temporarily inoperable, due to hardware or software failure, the nodes in the network continue to operate. Setting up the license server is discussed in further detail below. This section deals exclusively with configuring the License Manager, and assumes that a primary license server (and a secondary license server, if required) already exists and is running.

Once the licensing mode is changed to Client/Server Mode, a number of fields must be filled in. These fields are the primary IP address, secondary IP address, primary port number, and secondary port number. The primary IP address and primary port number are mandatory, and must be filled in with the IP address and port number of the primary license server. (The port number of the license server should be left as the default, 6797, unless a user is specifically

running the license server on a different port number). The secondary IP address and port are optional, and should be filled in with the IP address and port number of the secondary license server if available, and 0.0.0.0 port 6797 otherwise.

8. *Using the Provisioning System*

5 The provisioning system, described above, is used specifically for the management of data in the database servers. This section introduces the provisioning system, and provides basic information on how to use the provisioning system. The functionality offered by the provisioning system, and how to use this functionality; is described further below, when User Management and CDR Collection via the database server is discussed in detail.

10 The provisioning system is started in a similar way to the management system, using Start\Programs\Present Invention\Provisioning system. Before the provisioning system starts, it will require the selection of what type of records will be modified. When starting the provisioning system for the first time, simply select "User Records."

15 The first step in configuring the provisioning system to connect to a database server is to configure its identification information. This process is similar to the process used when setting up the management system. To access the identification configuration of the provisioning system, select "Security" from the "Connection" menu. A user will be presented with a security domain configuration dialog identical to the one presented by the management system. Use this dialog to enter information about a user's security domain in the same way as for the management system.

20 Once a user has configured the security domain settings for the provisioning system, the next task is to configure the list of database servers with which the provisioning system interfaces. The list of servers is available by selecting "List of Remote Servers Available" on the "Connection" menu. An Address Book 6602 is shown in FIG. 66 with a database servers - Address Book

panel 6604 which lists the remote servers available. The Current Server panel 6606 shows the name of the server currently selected.

Database servers can be added and removed from the list of available servers, using the Add and Remove button within the database servers - Address Book panel 6604. The IP address and port number is required for any database server that a user wishes to add. A user may also change the IP address and port number for database servers that have already been configured. Also, one needs to ensure that the identity of the provisioning system is in the client list on the database server, or it will not be allowed to connect.

With the provisioning system, one of the database servers available should be designated as the active server. The active server will be used to process all requests made through the provisioning system. If a user changes the active server, he or she should close any active session, and open a new session in order to perform operations. Unlike the management systems, the service should be started on the database server before the provisioning system will be able to connect.

Once a database server list and active database server have been configured, a user should be able to use the provisioning system to interact with data in the database. By selecting "New Session" from the file menu, and selecting the kind of session that one wishes to create, one can establish a connection to the database server, and can issue requests and manipulate data. How to use the provisioning system to do so depends, however, on the type of session and the operations that are being performed. Performing these operations are described in detail in later sections.

9. Troubleshooting Service Startup

This section provides summary information about the possible causes that can prevent a service from starting correctly. First, it is necessary to differentiate between not running, and not starting. If an operational component cannot run, then it will fail to load, and no icon will be displayed in the system tray. Possible causes for failing to run are:

- Insufficient access to configuration information (caused by user account being used to run the program);
- DLLs or other files installed with the software have been removed, or relocated;
- Another instance of the operational component is already running on the same machine.

Once a service is running, there are very few things that could possibly prevent it from starting. Failing to start is characterized by a red icon appearing in the system tray, and turning yellow while it is attempted to start the service, but returning to the red (stopped) state without turning green. The causes for a failure to start the service depend on the type of service.

a. Gateway Service

There are no known reasons why the Gateway Service would fail to start. Even if hardware devices fail to initialize, the Gateway Service should initialize nonetheless.

b. Routing Service

The Routing Service should start successfully in all cases. There are no known causes that can prevent the routing service from starting successfully.

c. Database Service

The Database Service may fail to start in the event that a connection to an ODBC database cannot be established. If this is the case, then the management system should be used to correct the ODBC connection information and the service startup should be retried.

II. OPERATION

1. Gateway Server Capabilities

This section provides an introduction to the gateway server, a component in an present invention network. In one example, the gateway server is the largest and most complex component of the present invention software, and is the component which brings together the functionality offers by various components of the system, and actually makes that functionality available to the end users of the system.

The material in this section is divided into several sections. A brief description of these sections is as follows.

Hardware Support. The present invention gateway server is designed to interact with a variety of underlying hardware, in order to support a diverse range of technology for interconnection to traditional telecommunications networks. Because of this, configuring hardware support in the software can be complex, and is explained in detail in this section. Normally, this process is not necessary, but information is provided for reference purposes.

Lines and Line Groups. Each line connecting the gateway server to the telecommunications network, as well as lines connecting the gateway server to other gateway servers, need to be configured. In addition, to simplify administration, lines are grouped into line groups, which allow certain properties

to be configured on a number of lines. This section explains lines and line groups, and how to perform configuration operations on them.

Services. At the heart of the gateway server is the services that it provides to the end users. The gateway server has an extremely flexible architecture that allows services to be seamlessly integrated with the system, added and removed, and routed between. An API even allows the development of custom applications that operate with the gateway server. This section describes how services work, and how to configure services for use.

Call Rating. The gateway server has the capability to debit for calls, explained in more detail in the services section. However, in order to charge for calls, the gateway server requires call rating information to be supplied, which govern how much calls cost. This section explains the call rating process, and how to configure call rates.

Voice Prompt Management. The gateway server features customizable voice prompts, with support for multiple languages. This section explains how to set up and customize voice prompts for use with the system, and provides some additional detail on how voice prompts can be administered.

Some of the capabilities that the gateway server is capable of providing require in the interaction of the routing server and/or database server. The use of these components to enhance the services that the gateway server provides is described in later sections.

1. Hardware Support

This section describes how the gateway server supports hardware devices, and how to configure the hardware support component of the gateway server

In order to connect a gateway server to a traditional telecommunications network, specialized hardware boards must be used. The gateway server supports a range of such boards, each offering different numbers of lines, and types of lines. In addition to these telephony boards, the gateway server is also capable

of making use of DSP processing boards that enable larger numbers of calls to be handled by a single gateway server. The hardware support component of the gateway server is responsible for interfacing with the installed hardware devices, and on enabling those devices to co-operate in providing IP Telephony services.

5 In addition to supporting telephony devices, the gateway server extends the concept to network devices, used for interconnection to IP Telephony devices via an IP network. In one example, the only physical device installed in the gateway server for this interconnection is an Ethernet card.

10 In addition to supporting a variety of types of hardware devices, the gateway server can also support multiple devices of a single type. For example, it is possible to use two Dialogic D/41ESCs, each of which is a 4 line telephony boards, to achieve 8 lines of capacity.

15 A Hardware Device Configuration screen 6702, as shown in FIG. 67, can be accessed by double clicking on Hardware Devices, a top-level management tree node on the gateway server. This window 6704 will display a list of devices that is currently configured, and the current status of those devices.

20 From this screen in FIG. 67, individual hardware devices can be started or stopped. By default, when the gateway service is started, all hardware devices will be started, and when the gateway service is stopped, all hardware devices will be stopped. Add and Remove buttons can be used to add and remove hardware devices, and the Properties button can be used to configure existing hardware devices. When a hardware device is added. Since its properties are configured, a type-specific dialog will appear that allows a user to configure information for that specific hardware device.

25 The subsections below discuss each hardware device supported by the system. When one chooses to add a hardware device, he or she will be prompted to select what kind of hardware device is being added from one of the following devices. Some devices allow multiple instances to be present. Because of this, devices are identified by the combination of their type, and an instance number

assigned to the device when it is added. Prior to discussing each individual board, however, general information regarding Gatenet 3 vs. DSP modes of operation are provided.

a. Gatenet 3 vs. DSP Compression

5 The gateway server uses complex signal processing algorithms in order to minimize the bandwidth required for communication between Gateways, while simultaneously preserving a high quality of voice. There are two ways that the required signal processing can be performed:

10 (1) Echo cancellation, the first step in processing voice data, is performed on the hardware device that is connected to the telephone equipment. However, the CPU of the machine used to run the gateway server is used to apply compression algorithms to the voice data & known as a Gatenet 3 configuration, allows the use of a single (or multiple) telephony interface boards without a requirement for a specialized DSP processing board. The drawback to this approach is that there is a limit to the number of simultaneous calls that can be handled by a main CPU of certain types of computer that run the gateway server. A second drawback is that some devices are reduced in terms of the number of lines they can support when echo cancellation is performed onboard. In some cases, not intended to limit the invention, only certain boards can be
15 operated in Gatenet 3 mode;

20 (2) Likewise, echo cancellation and voice compression are performed on a specialized DSP board that is connected in hardware to the telephony interface board. This provides better performance and scalability, but incurs the cost of the specialized DSP board.

25 The gateway server software is capable of supporting either mode. However, since the hardware configuration depends on which mode is selected, the mode of operation is normally predetermined.

b. Dialogic D/21H

The Dialogic D/21H operates only in Gatenet 3 mode, but is completely compatible with the Dialogic D/41ESC when operated in this mode. Please see the section below on the Dialogic D/41ESC for details on how to configure a Dialogic D/21H.

c. Dialogic D/41ESC

The Dialogic D/41ESC is a four port, analog loop start interface board. It allows a connection to a central office using standard analog telephone lines, as well as interconnection with PBXs or other equipment that provide FXO (also called IST) ports.

The Dialogic D/41ESC is capable of operating either in Gatenet 3 mode, or in DSP mode. When operating in Gatenet 3 mode, it is only capable of providing two ports of connectivity to telephone equipment; otherwise, all four ports can be used. Also, a D/21H functioning in Gatenet 3 mode is identical to a D/41ESC, and can be set up and treated in exactly the same manner.

d. D/41ESC General Panel

The first panel displayed is a General panel 6802 as shown in FIG. 68. Other panels can be selected by clicking on the Tabs 6804 that appear along the top of the dialog. The settings on the General panel are explained below.

Board Number. This is the Dialogic board number of Dialogic D/41ESC being configured, and controls which board (if more than one D/41ESC is present) will be used. Normally, the first D/41ESC device added should be assigned board number 1, the second 2, and so forth. However, if some boards in the chassis should not be used, or are used by other applications, this parameter allows explicit specification of which Dialogic board number to open.

Enable GN3/EC Mode. This parameter is used to enable or disable operation in Gatenet 3 mode. See the section on Gatenet 3 vs. DSP compression above to determine how this parameter should be configured.

5 *Number of rings before answering.* This controls how many times an inbound call will ring before the gateway server answers the call.

10 *Enable Trunk Channels.* This group of buttons allows a user to select specifically which trunk channels will be opened on the D/41ESC. Note that if one enables Gatenet 3 mode, one cannot enable trunk channels 3 and 4. Normally, one should enable all channels unless using other channels of a particular Dialogic board with a different application.

15 *Enable Voice Channels.* This group of buttons is the same as the trunk buttons, but control which voice channels are used. Under most circumstances, one uses the same settings for voice channels and trunk channels.

20 *Output Gain.* This setting allows one to control the output gain. If the volume of voice prompts and speech are too low when making calls, then adjusting this option allows the gain to be adjusted. However, before using this parameter, one should verify that other Dialogic options are set correctly, since gain can be applied in several different locations. Also, if one uses a separate DSP board for processing, this setting will have no effect since the DSP board will apply gain, not the D/41ESC.

e. D/41ESC Call Analysis Panel

A call analysis panel 6902 is used to configure how the D/41ESC detects the presence of a connection. This call analysis panel 6902 is shown in FIG. 69. It is a panel in window 6802 and continues to include Tabs 6804.

25 The parameters on the call analysis panel are passed directly to the corresponding Dialogic modules that make use of these parameters. For a detailed discussion of how the Dialogic D/41ESC performs call analysis, refer to the Dialogic documentation. The two most relevant parameters on this page are

Enable Voice Detection, and Enable Loop Current Signaling. Voice detection should always be enabled, except in specialized environments. Voice detection allows the beginning of a call to be determined based on the recognition of speech patterns, such as “Hello”. Loop current signaling allows loop current based answer supervision provided by the telephone network to be interpreted, and should only be enabled if supported by the telephone network provider.

f. D/41ESC Ring/Dialtone Settings Panel

The D/41ESC can be configured to wait for a dial tone before it enters digits, as shown in panel 7002 in FIG. 70. It can also be configured to wait a maximum number of rings on an outbound call before it gives up, and assumes that the call will not be answered. This panel allows one to configure settings that affect this process.

Ring Settings: Time for Outbound Call Pickup. This parameter controls how long the Dialogic D/41ESC will wait for an outbound call to be picked up before it abandons the call.

Dial Tone Settings: Minimum Dial Tone Length. This parameter controls the minimum length of a dial tone before it will be recognized.

Dial Tone Settings: Timeout for Dial Tone. This parameter controls how long the Dialogic D/41ESC will wait for a valid dial tone before it abandons the call due to no dial tone.

Dial Tone Settings: Maximum Dial Tone Interruption. This parameter controls the maximum break in a dial tone that will be tolerated without considering a particular tone not to be a valid dial tone.

g. D/41ESC Tone Definition Panel

During Call Progress Analysis, and also for hang-up detection, the D/41ESC uses cadence detection to determine when a line is ringing, when a user answers, when the line is busy, and so forth. However, many of these tones may not be common from one country to another. As a result, it is possible to customize the tones that are used by the Dialogic D/41ESC in order to determine whether or not particular events are happening. The tone definition panel 7102 is shown in FIG. 71.

By selecting one of the tones listed, and clicking on Edit, one is able to modify the frequencies and timings that control how a particular tone is detected. One should determine the information that defines the tones that are used on the particular equipment that one is connecting to the gateway server. When editing each tone, a cadence dialog 7202 is displayed which asks for settings as shown in FIG. 72.

One should supply the information in the dialog. The parameters are fairly self-explanatory. By default, dual frequency cadence tones will be detected. If one wishes to detect a single frequency tone, then set the second frequency to zero. To detect a continuous tone, not a cadence tone, one should set the repeat count to zero.

The two most important tones in most cases are the hang-up cadence tones. Hang-up cadence tones are tone definitions that can trigger the automatic disconnection of a call. In some environments, disconnection signaling is only performed by way of cadence tones. In this case, it is necessary to enter the appropriate tone information under either the primary or alternate hang-up cadence. It is also necessary to enable disconnection upon detecting this tone at the line level.

h. Dialogic D/160SC-LS

The Dialogic D/160SC-LS is very similar to the D/41ESC, but with two notable exceptions:

- It has 16 analog loop start interfaces instead of 4;
- It cannot operate in Gatenet 3 mode.

This board is suitable for customers who require a large number of analog lines connected to a telephone device. The configuration panels for the D/160SC-LS are identical to those for the D/41ESC, except for general panel 7302, which provides a larger number of ports that may be enabled or disabled. The D/160SC-LS panel is shown in FIG. 73.

For all other panels 7304, refer to the information provided for the D/41ESC above.

i. Dialogic D/240SC-T1 (Robbed Bit T-1)

The Dialogic D/240SC-T1 is a T-1 interface board that provides 24 channels of digital connectivity to a telephone device with a T-1 interface. Although the Dialogic D/240SC-T1 is capable of operating either in ISDN PRI mode, or in robbed bit mode, the Dialogic D/240SC-T1 device in the gateway server refers only to the D/240SC-T1 operating in robbed bit mode. If a user is operating the D/240SC-T1 in ISDN PRI mode, one should add a Dialogic ISDN device, described below.

Much of the configuration of the D/240SC-T1 is identical to the configuration for the D/41ESC. All of the T-1 specific configuration is contained on general panel 7402 as shown in FIG. 74. Tabs 7404 contain identical information to the configuration or the D/41ESC.

The configuration fields on this page are as follows:

DTI Board Number. A D/240SC-T1, like a D/41ESC, provides both T-1 trunk and voice resources. However, there are separate logical board numbers for

the trunk and voice boards. This parameter allows one to control which DTI board number will be opened for this device.

Voice Board Number. As explained above, the Dialogic voice board number for a D/240SC-T1 is opened separately from the T-1 trunk. The D/240SC-T1 actually emulates 6 independent voice boards. Thus, for D/240SC-T1 # 1, the voice board number starts at 1. For D/240SC-T1 # 2, the voice board number starts at 7. Normally, one will not need to customize settings made by ExpressConfig, however, for integration or sharing with other applications, this option is provided.

T-1 Signaling Bits. With T-1 robbed bit protocols, information is signaled between the telephone equipment and the gateway server using two bits, known as the A and B bits. Different patterns for these bits represent different information. Most telephone equipment uses the default settings for the T-1 signaling bits. However, the interpretation of bits received from the telephone equipment, and the bits that will be sent to the telephone equipment, can be customized in this page.

T-1 Network Mode. This mode controls how the T-1 signaling bits, defined above, will be used. It controls the relative timing and sequence of the bit patterns. One should select the option that corresponds to your telephone equipment.

j. Dialogic ISDN

The Dialogic ISDN device provides support for both T-1 ISDN PRI systems and E-1 ISDN PRI systems. If one has configured either a D/240SC-T1 or a D/300SC-E1 in ISDN PRI mode, then one should install a Dialogic ISDN device. One Dialogic ISDN device is required for each ISDN span (23 lines for T-1, 30 lines for E-1) that one has configured in a system.

The configuration of the Dialogic ISDN device is significantly simpler and more understandable than dealing with the cadence tones and settings for

some of the other Dialogic devices. This is because with ISDN, signaling is explicit, out of band, and relatively constant between different environments. There are three panels that a user deals with when configuring a Dialogic ISDN device.

(1) General Panel

General panel 7506 of properties screen 7502 allows general settings to be configured on the D/240SC-T1. FIG. 75 shows this panel. Other tabs 7504 are shown in FIG. 75 and discussed in detail below. These settings are explained below.

DTI Board Number / Voice Board Number. These parameters are identical to the settings explained for the D/240SC-T1. Refer to that section above for an explanation of these parameters.

Board Type. This parameter controls whether the ISDN device in use is a D/240SC-T1 or a D/300SC-E1. You should select the parameter corresponding to the board that is installed in your system.

Output Gain. Dialogic ISDN devices are almost always used with a DSP voice board, so the output gain parameter should not normally be used unless one is involved in a very specialized application.

(2) Channel Configuration Panel

Channel configuration panel 7602 allows one to select which channels will be used on the Dialogic ISDN device. In almost all cases, one should disable all voice channels, since the voice channels are not used. Although enabling the voice channels does not interfere with the gateway server, some ISDN protocols require enough DSP processing power that the voice channels cannot be enabled successfully. In most cases, one should enable all of the trunk channels. FIG. 76 shows the channel configuration panel.

(3) Call Setup Parameters Panel

Unlike conventional telephone devices, which signal calls using the hookswitch state and DTMF tones, ISDN devices signal calls in a single digital message. The contents of this message do vary between ISDN switches, however. Call setup panel 7702 allows the correct message for a particular ISDN switch to be sent. FIG. 77 shows the call setup parameters panel.

The parameters in this panel are as follows:

DestinationNumber Type. Indicates the type of number being transmitted to the ISDN switch. This field needs to match the type required by the ISDN switch one is connected to.

Destination Number Plan. Indicates the numbering plan used for the destination number transmitted to the ISDN switch. This field needs to match the numbering plan required by the ISDN switch being connected to.

Connection Timeout. This controls how long the gateway server will wait for the completion of a particular call. If the call is not completed within this amount of time, it will be aborted.

Originating Number Type. Same as the destination number type, but for the originating number (i.e. the number of the gateway server).

Originating Number Plan. Same as the destination number plan, but for the originating number.

Originating Number. Number sent to the ISDN switch as the number of the party originating the call. Normally, when connected to the PSTN, this parameter will be ignored, since the PSTN switch knows the identity of the originator.

Layer 1 Protocol. Controls the value of the Layer 1 Protocol parameter in the call setup message. This should normally be left at the default value, but some specific ISDN switches may require this value to be overridden.

k. Dialogic MSI/SC

The Dialogic MSI/SC boards are analog station interface boards. The MSI/SC allows telephones, or other equipment that normally plugs into a standard analog line from the PSTN, to be connected directly to a gateway server. In essence, the MSI/SC emulates a line from the PSTN, and can be used to connect to PBX trunks in addition to standard telephones. This kind of interface is also known as an FXS interface, or station side interface.

There are two variants of Dialogic MSI/SC boards, those with the ringing option and those without. In order to use the Dialogic MSI/SC with the gateway server, the ringing option is used in order to make outbound calls.

In addition, there are three configurations of an MSI/SC, an 8 line, 16 line, and 24 line configuration. These are referred to as the MSI/80SC-R, MSI/160SC-R, and MSI/240SC-R, where the -R denotes the ringing option. A user needs to know which of these configurations he or she has prior to configuring the MSI/SC for use with the gateway server.

Panel 7802 is configured for an MSI/SC board as shown in FIG. 78. The fields in the panel are explained below.

Board Number. This controls which Dialogic MSI board number will be opened.

Board Type. Indicates the type of board, either the MSI/80SC-R, MSI/160SC-R, or MSI/240SC-R.

Ring Line For Outbound Calls. Indicates whether or not the ringing option is present on the MSI/SC board. This option is used in order to make outbound calls.

Unanswered Rings before Aborting. When ringing a particular line, the MSI/SC will generate a certain number of rings before the call is aborted. This parameter allows the number of rings that will be generated before an abort to be configured.

Output Gain. This parameter allows the gain of the lines connected to the MSI/SC to be controlled. This can be used to adjust the volume of the phones connected to the MSI/SC.

ANI Information. ANI, or Caller ID information, can be used by gateway server applications in order to identify a user without asking for identification. Since the lines connected to the MSI/SC may be fixed (i.e. line # 3 is always connected to the phone of a particular user), Caller ID information can be associated with each port of the MSI/SC. One can use this control to enter Caller ID information for each MSI/SC line.

1. *Mibridge DSP*

As discussed in a section above, the gateway server supports the use of specialized DSP processing boards to achieve high density and consistent call quality. Although boards from different manufacturers are supported, the firmware used on the boards is licensed from Mibridge, Inc. in all cases. The Mibridge DSP device in the gateway server refers to any of the DSP boards supported by the Mibridge firmware, and any quantity of those devices. A user should install only one Mibridge DSP device, even if he or she has multiple DSP boards.

There are two simple panels that should be configured for the Mibridge DSP device. The Board Type panel 7902 has only one option, allowing the type of board to be selected. One should select the board type that corresponds to the actual hardware present in the gateway server. A voice channel configuration panel 7904, shown in FIG. 79, allows addition options to be configured. The options are explained below.

SCbus PCM Encoding. Depending on the Dialogic devices being used with the Mibridge DSP board, the PCM encoding on the SCbus used to connect the Dialogic and DSP boards may be different. If one is using any board except

for the D/300SC-E1, mu-Law encoding is normally used. If, however, one is using the D/300SC-E1, one should select A-Law encoding.

5 *Automatic Gain Control.* In certain environments, it is beneficial to enable automatic gain control to compensate for naturally low input levels. This is controlled via the Automatic Gain Control setting. In most environments, however, Automatic Gain Control should be disabled.

10 *Fixed Input/Output Gain.* If automatic gain control is not enabled, then a fixed level of gain can be configured. In most environments, the default gain of 0db should be used. However, the level of gain can be customized if the default setting is not sufficient.

m. Network H.323

15 The Network H.323 device enables communication using the H.323 protocol. This allows calls to be received from and placed to H.323 compatible devices, such as Microsoft NetMeeting terminals or multimedia PCs. The only required configuration parameter for the Network H.323 device is the H.323 terminal name that will identify the gateway server to the remote H.323 entities with which it communicates.

n. Network Proprietary

20 The Network Proprietary device enables communication with other present invention gateway servers using a proprietary protocol. This is the normal protocol used for Gateway to Gateway communication. There are two main configuration parameters for the Network Proprietary device. The first is the port range, which controls which UDP/IP ports will be used for media data transmitted to and from this gateway server. The second is the Local System ID
25 of the gateway server. This name will be presented to remote gateway servers when a call is placed, and may be used by the remote gateway server in order to

identify and authenticate the originating gateway server. More details on this process are discussed in the User Management section.

The “Network Settings” tab allows advanced network properties to be configured and tuned. These settings, as well as a detailed explanation of the port range feature, are found in Appendix B, which provides technical network design information.

2. *Lines and Line Groups*

The hardware support component of the gateway server, discussed in the previous section, provides support for hardware devices that are located in the PC on which the gateway server runs. However, providing actual functionality to the end user involves co-operation between different boards. Furthermore, lines on different boards may often need to be treated the same way, and pooled together. The line and line group component of the gateway server organizes the capabilities provided by underlying hardware devices into virtual “lines”, and allows these lines to be grouped together into line groups. This section explains lines and line groups.

a. *Lines*

Lines represent a single logical point of connection to the system. A line is capable of handling a single call from a user, and represents one complete port of connectivity. There are two kinds of lines; telephony lines, which are directly related to the telephone lines actually connected to the gateway server, and network lines, which are created on demand and used to represent connections to other gateway servers or to users on the H.323 network.

Since network lines are created on demand, it is not necessary to interact with network lines in most cases. However, telephony lines are created in advance, and must be created in order for the gateway server to make use of the underlying hardware devices. Telephony lines serve primarily to combine

resources supplied by underlying hardware devices in order to deliver the required services to the user. For example, a DSP device is not capable of actually connecting to the PSTN, but it is capable of performing IVR. On the other hand, a telephony board may not be capable of performing IVR, but can manage the connection to a user. Telephone lines group these resources together to form a single virtual line that can be used to deliver the services provided by the gateway server.

The actual configuration of a telephony line can be quite complex, and is in general not necessary since ExpressConfig will set up suitable default parameters. The configuration of telephony lines is discussed below when the telephony line group is discussed in detail.

Lines are either in the active state, the inactive state, or switching between the two states. Network lines are always in the active state, since they are deleted once they are no longer active. Lines can be activated or deactivated at any time. When the gateway service starts, all lines will automatically be activated. When the gateway service stops, all lines will be automatically deactivated.

b. Line Groups

For the purposes of call routing, as well as configuration of parameters, lines are always organized into groups of lines, called line groups. Line groups manage properties that generally do not need to be configured on individual lines. Also, the service level of the gateway server may associate application-specific properties with line groups that affect how services are provided to the user. For example, the service layer could be configured to handle all calls coming in on a particular line group by routing the call to a specific number, whereas calls to another line group would be handled using an IVR application.

As with lines, there are two kinds of line groups; telephony line groups, and network line groups. Telephony line groups contain numbers of telephony lines. Network line groups contain network lines, and handle the dynamic adding

and removal of network lines as calls are established and terminated. The configuration of these two types of line groups differs; each will be discussed in a separate section below.

Using the management system, a line groups window 8002 can be displayed by expanding the top level configuration tree node named "Line Groups". This window is shown in FIG. 80. In the top part of the window, a line group panel 8004 shows that line groups may be selected from one of the available line groups. Typical installations have one network line group, and one telephony line group. Once a particular line group is selected, the bottom part of the window displays a telephony line group status panel 8006. Panel 8006 shows the list of lines that are currently configured on that line group. The state of each of the lines will also be displayed, along with any error messages associated with that line. In the top window, line groups can be added, removed, or configured. In the bottom window, individual lines can be activated or deactivated, and lines can be added, removed, and configured. The configuration of lines and line groups are discussed in the sections below.

c. Telephony Lines and Line Groups

Telephony line groups manage groups of telephony lines, and apply parameters and properties that are common across all lines in that group. Telephony line groups also interpret events that occur on lines, such as DTMF digits, and can assign a meaning to those events. For example, a hardware device or line might only be aware that the DTMF digit sequence '###' was pressed; however, the telephony line group may know that this sequence indicates that the call should be disconnected.

When a telephony line group is added or configured, dialog 8102 with general panel 8106 appears and allows parameters to be customized. Panels can be navigated between using tabs 8104 at the top of the dialog 8102. The settings

available on the general configuration dialog are shown in FIG. 81 and described below.

Maximum Connect Time. If one wishes to limit the maximum duration of a connection to the system, it is accomplished using this parameter. Normally, one will not need to enable a maximum connect time, but depending on the application, it may be necessary to do so.

Line Allocation Strategy. The line allocation strategy controls how a line is selected when an outbound call is to be placed using the telephony interface. These strategies are usually used in conjunction with the line allocation strategy used by the telephone equipment to which the gateway server is connected in order to minimize glare conditions. There are four possible strategies:

- Lowest line allocated first. The lowest free line number will be used to place the call;

- Highest line allocated first. The highest free line number will be used to place the call;

- Round robin (decreasing). The first call will be placed with the highest line number in the system. Each successive call will be placed using the highest available line number that is below the previous line number used to place a call. Once no lines are available, the scheme starts again with the highest available line number;

- Round robin (increasing). Identical to round robin (decreasing), but starts with the first line and progresses up.

Enable Support for DID. Enables the use of DID information on the lines in this line group. This should only be enabled if the telephone equipment to which the gateway server is connected provides DID signaling.

DID is sometimes used to refer to DNIS service provided on an ISDN PRI line. One should not enable DID on an ISDN PRI line. Enabling DID only allows the collection of DID information using DTMF digits sent at the beginning

of a connection. ISDN PRI does not signal DNIS (sometimes called DID) information in this fashion.

DID Min/Max Digits Accepted. If DID is enabled, controls how many DID digits will be required in order for DID collection to be successful. If the number of digits is between the maximum and minimum, inclusive, then the DID information will be accepted; otherwise, DID collection will fail.

DID Collection Time Limit. Specifies how long DID digits will be collected for after the call is answered. Normally, one second is sufficient for the transfer of DID digits, but this may depend on your particular telephone service.

(1) Hangup Settings Panel

Hangup settings panel 8202, shown in FIG. 82, allows the events that cause a user to be disconnected to be configured. In normal environments, it is not necessary to configure these settings. However, in some environments, disconnection is not signaled explicitly by the telephone equipment, and must be detected by alternate means.

Silence Triggered Hangup Settings. These settings allow one to configure whether or not a line will automatically be disconnected based on an extended period of silence. Normally, this feature is disabled. When enabling the feature, one should select whether silence of the local party, or silence on both ends of the conversation is required to cause disconnection. One should also specify the time, in seconds, of continuous silence that is required in order to cause disconnection.

DTMF Triggered Hangup Settings. When integrating with PBXs and other kinds of telephone devices, disconnection can sometimes be signaled through the use of DTMF digits. If one enables DTMF triggered hangup, one will be required to specify a string of DTMF digits (including extended digits such as A, B, C, D) that will cause hangup. If these digits are detected in the

order specified, and within the amount of time specified in the interval field, then the call in progress will be disconnected.

(2) Integration Digit Settings

The integration digit settings are used for custom applications. One should not modify these settings.

The above describes the settings for the telephony line group. In addition, when configuring a line on the line group, it is necessary to provide configuration information for that line. In order to understand telephony lines, it is necessary to understand how resources are represented in the system. There are three types of resources associated with a telephony line. The trunk resource represents the actual connection to the phone equipment, such as the physical analog line, or a T-1 channel, and so forth. The voice resource represents the DSP processing resource that is used to perform IVR, voice compression, and so forth. Finally, an optional auxiliary resource can be used for call progress analysis, cadence detection, and other such applications. ExpressConfig normally sets up the resources for each telephony line exactly as required. It is generally not recommended that one modify these settings.

When a telephony line is added or configured, the following set of panels can be interacted with in order to configure the telephony line.

(3) Resources Panel

The resources panel 8306 controls the interaction of the telephony line with the underlying hardware devices. Tabs 8304 allow selection of other panels which are discussed in detail below. Panel 8306 is configurable from telephony line window 8302, which is shown in FIG. 83, and explained as follows

Location/Position Information. This field allows one to provide identifying information that can enable you to locate the physical line. It is for descriptive purposes only, and is not required.

Trunk Resource Configuration. In this section, one needs to specify the hardware device that is associated with this telephony line, and the instance of the trunk resource on that device which is to be used. Trunk resources are described above.

5 *Voice Resource Configuration.* This section allows one to specify which voice resource should be associated with the trunk resource configured above. Voice resources are described above.

(4) General Panel

10 The parameters provided for general panel 8402, shown in FIG. 84, follow:

15 *Accept Any Location/Location List.* These settings are for future expansion, and will allow one to address individual lines when routing calls. This enables the development of virtual PBX style applications, but is not required for operation with any of the applications currently shipping with the gateway server.

20 *Class of Service.* The class of service on each line can be configured on a line-by-line basis. Class of service constraints allow lines to be disabled, to accept only incoming calls, to place only outgoing calls, or to accept calls in both directions (duplex service). The desired class of service can be specified using this setting.

d. Call Analysis

This panel allows you control call analysis settings (panel not shown, but tabs 8304 show tab to access this panel). There are two main settings:

25 *Enable Call Progress Analysis.* This parameter determines if call progress analysis should be enabled. If enabled, it is also necessary to select which resource will perform the call progress analysis. The correct value for this

setting is generally determined by ExpressConfig, and depends on a particular hardware configuration.

Enable Cadence Detection. This parameter determines if cadence detection is active during calls. Cadence detection may be used to trigger certain things, such as automatic disconnection. If cadence detection is enabled, it is necessary to specify which resource will perform the cadence detection.

Auxiliary Resource. If either cadence detection or call progress analysis requires the use of an auxiliary resource, the auxiliary resource must also be configured. It is necessary to specify the device and instance of the resource associated with this line, just as for trunk and voice resources.

(1) Connection Detection Panel

Panel 8502, shown in FIG. 85, controls how the start of a conversation is detected by the system. There are several options for how this connection detection occurs.

Connect Immediate. If this parameter is enabled, all calls are assumed to be connected after they have been dialed, without waiting for any signaling from the telephone equipment.

Connect CPR. If enabled, this parameter will enable calls to be connected if call progress results are provided by the call analysis resource associated with this line.

Connect LCR. If enabled, this parameter will allow loop current reversal events to trigger the connection of a call. In some environments, loop current reversal is used for answer supervision.

Connect Cadence 1. Enables connections to be signaled through the use of custom cadence tone 1, which must be defined at the hardware device level. If a call is being placed, cadence tone 1 is detected, and this parameter is enabled, the call will be connected immediately. Note that the process of placing a call

may interfere with the use of cadence detection in certain hardware configurations.

Connect Cadence 2. Identical to the above, but for custom cadence 2.

(2) Disconnection Detection Panel

5 The disconnection detection panel (panel not shown, but tabs 8304 show tab to access this panel) is very similar to the connection detection panel. However, events in this panel are used to detect the disconnection of the user previously connected on a particular telephony line.

10 *Disconnect LCR.* Enables or disables loop current reversal signaling of disconnection.

Disconnect LCOff. Enables or disables loop current off signaling of disconnection.

Disconnect Cadence 1. Enables or disables disconnection signaling using custom cadence tone 1.

15 *Disconnect Cadence 2.* Enables or disables disconnection signaling using custom cadence tone 2.

(3) Network Lines and Line Groups

 Unlike telephony lines, network lines are created dynamically by the system as new calls are received, and are removed when calls are terminated.
20 Therefore, there is no need to configure a list of lines, or to set configuration for those lines.

 Furthermore, the network line group itself has no configuration properties at the line group level. (Recall that settings may be associated with line groups at the service level, described in the next section). The reason for this is that
25 parameters are associated with the network proprietary device, or the network H.323 device.

A second network line group is generally not required. All network calls are processed through the single network line group. However, one must have a network line group in order to receive network calls.

3. *Services*

5 The Services component of the gateway server is shown in FIG. 3B. This component is responsible for enabling the applications that actually provide services to the end user.

Applications are highly separated from the gateway server platform 304 discussed so far in this section. Applications make use of the numerous capabilities provided by the gateway server platform 304 in order to interact with the user, and to deliver enhanced services. In one example, all such applications are written using a C++ API that provides access to functionality provided by the gateway server. Services can interact with each other, and transfer control from one service to another, creating a flexible environment for constructing applications.

10 Gateway server applications are contained in DLL files that are linked at run-time to the gateway server platform 304 and operate in gateway server Application Layer 302. However, it is possible for third parties to create similar application DLLs. It is also possible to use the gateway server API to write applications that run from a separate process, or even a different machine.

15 The first part of this section will describe the services component of the gateway server platform 304. The second part of this section describes the applications according to the present invention that are included with the gateway server.

20 Selecting the "Services" node in the configuration tree using the management system will access a main service window 8602 shown in FIG. 86. The area at the bottom of the window shows a list window 8604 of active sessions, the line that each session is associated with, and the service/application

that is currently managing that session. The other information presented on this screen is discussed below.

a. Service Selection

5 The gateway server is capable of supporting numerous kinds of services, and can run any number of instances of a particular service. Because multiple services are present that can handle an incoming call, it is necessary to define rules for selecting the service that will be used to handle incoming calls. By default, ExpressConfig creates a rule that selects the gateway application to handle all incoming calls. However, this behavior can be customized, and different services can be selected.

10 In order to select the service used to handle incoming calls, the gateway server provides a rule tree 334 within Rule Engine 332. When an incoming call is received, the gateway server will traverse rule tree 334 until it finds a service that matches the parameters of the incoming call. As previously discussed, calls are handled as sessions by the gateway server. Thus, incoming calls are a type of session within the system. Information in the incoming call can be used to control the service selection process. For example, DID or DNIS information can be used to select which service is activated, and with what parameters. This is extremely powerful in creating a flexible, single box solution that hosts multiple applications. FIG. 86 shows service selection panel 8608 and rule list panel 8606. Rule list panel 8606 allows rules to be added, modified and/or removed. This feature is discussed below.

15 The rule tree 334 consists of a number of rules, arranged in a tree like structure. When an incoming call arrives at that gateway server, selection begins at the top of the tree. Each rule at the top level of the tree is processed until a rule is found which matches the information contained in the incoming call. If a match is made, then that node in the tree is selected. A node in the tree, however, can either be a final node, in which case it contains the name of a

service used to handle the call, or it can be a subtree. If it is a subtree, then the above process is repeated starting at the top of the tree that is beneath the selected node. If the bottom of rule tree 334 is reached and a matching rule is not found, then no service will be selected and the incoming call will be rejected

5 Each rule in the rule tree 334 contains a condition that must be matched for the rule to be selected. A blank rule allows all incoming calls to be matched. If the rule is a final rule, meaning that it directly selects a service, then the name of the service, and parameters for the invocation of the service, are also required. Dialog 8702 opens to show settings for any rule. Dialog includes fields as shown in FIG. 87 and described below:

10 Rule Settings panel 8704 contains:

Attribute. This field provides the name of the parameter that will be checked to see if this rule should be selected. It names an attribute of the incoming call, such as DNIS, DID, or others.

15 *Relation.* This selects the relation that the attribute must have to the value, described next, in order for this rule to be selected. Possible relations are equals to, and not equals to.

Value. This is the value against which the value of an attribute will be compared. For example, if an incoming call has DNIS=1234, and the Attribute name is DNIS, then 1234 will be compared against the value specified in this field. In the example above, if the relation was equals to, and the value is 1234, then this rule will be selected. If, however, the value was 1235, or the relation was not equals to, then the rule would not be selected.

20 *Service.* If configuring a final rule, the name of the service that will be used to handle the call must be specified in this field

25 Bindings panel 8706 contains:

Bindings. In order to customize the behavior of a service, it is possible to specify additional bindings to be passed along to the service. Bindings are a list of (attribute,value) pairs that will be interpreted by the service that is invoked.

Clearly, the bindings list is service dependent. For a list of bindings that can be used to modify the behavior of a service, one needs to refer to the documentation for that service. As an example, placing "Language=English" as a binding will cause the gateway server to use English as the initial language for voice prompts. This is extremely powerful when combined with rules, because it would allow a configuration where calls received with DNIS=1234 are handled using English voice prompts, and calls received with DNIS=1235 are handled using Chinese voice prompts.

b. Service Level Line Group Settings

There are certain parameters that affect services and service selection that are required to differ from one line group to another. For example, it might be desirable to allow users calling in on one line group (connected to a PBX, for example) to make calls without entering an identification code, whereas calls received on a different line group might require an identification code. Service level settings can be associated with any of the configured line groups. These settings can be viewed by double clicking on the line group nodes in the configuration tree that appear when the service node is expanded. There is one node for each line group that exists in the system. The settings, which are common for all line groups, are shown in FIG. 88 and explained below.

User Lookup panel 8804 contains:

Auto-user Lookup using Call Setup Information. Information may be associated with incoming calls that allow a user to be identified automatically, without the use of user identification entry. For example, a user could be identified using Caller ID information supplied by the telephone equipment. Enabling this field will cause an automatic user lookup to be initiated as soon as the call is received, using all information available at the time that the call is received. More details on how user records are managed can be found under the User Management section of this document.

Auto-user Assignment using Call Setup Information. In the event that the lookup of a user based on Caller ID or other properties fails, or if this feature is not used, it is possible to always assign users to a particular user account. Because this is configured on a line group by line group basis, it could be used to assign all callers on a particular line group (e.g. one connected to a PBX) to a particular user account, perhaps PBX_USERS. On the other hand, calls on other line groups would continue to require user identification, whereas the users on the PBX interface would be automatically associated with the PBX_USERS account. If this field is selected, it is necessary to enter the user ID with which the line group will be associated, under User ID.

Default User Identification panel 8806 contains:

Default Privilege Level. Routes are restricted on the basis of privilege level, as will be explained in the section on number translation and call routing. Before a user is identified, however, they will have a certain privilege level. This parameter can be used to control the default privilege level assigned to an unidentified user. This in turn can be used to restrict calling privileges. For example, one route to 1-800-NEWUSER could be assigned a privilege 0, with all other routes requiring privilege 1. A greeting could say "If you are a new user, please call 1-800-NEWUSER to open an account. Otherwise, please enter the number you wish to call". By granting a default privilege level sufficient to make the call to 1-800-NEWUSER, it is possible for new users to call in and open accounts using the same system, without requiring multiple telephone lines and telephone numbers. Similarly, unidentified users on different line groups could have different default privilege levels.

DNIS/DID information panel 8808 contains:

Use DNIS/DID Information as Dialed Number. Enabling this parameter causes the gateway application, discussed below, to treat DNIS/DID information as the number that the user wants to call. This can be used to instantly route calls based on the number dialed by the user to access the system. This can be used

to provide 800 services, for example, by routing calls to 1-800-1234567 to customer A's local number of (703)555-1234, and routing calls to 1-800-7654321 to customer B's local number of (416)555-1234. This allows a single line to be used to provide 800 services to multiple customers that are located in different cities, reducing their long distance costs for handling calls from customers.

c. Gateway Service

The gateway service, in the context of high level applications, is a core service that leverages the capabilities of the gateway server platform 304 to deliver calling capabilities to end users. The gateway service is provided by the gateway server. It includes, but is not limited to, handling calls, bridging communications systems, and providing access to other systems, such as databases, routing servers, etc. The gateway application service should not be confused with the gateway server service. Because the Gateway Application Layer 302 is complex in its operation and provides a significant degree of functionality, it is discussed in the next section.

d. Autoforward Services

The autoforward services 313 is a relatively simple service that forwards all calls to a pre-determined number. It is capable of playing a greeting immediately before the call is forwarded. This service can be useful in a number of applications. For example, a particular DNIS field could select the autoforward service, which could announce "Please hold, and the next available customer service representative will be pleased to serve you", and automatically forward the call to the customer service department ACD – which could potentially be in a different geographic location.

Autoforward services 313 can be configured by double clicking on the "Autoforward" node in the configuration tree, which can be accessed by expanding the "Services" root level node. An Autoforward configuration dialog

8902 is shown in FIG. 89 and described below. Note that one must click on “Change” in FIG. 88 order to confirm changes to the service settings in call address panel 8904 and voice prompt panel 8906. These panels are discussed below.

5 *Country Code/Area Codes/Local Number.* These fields specify the number that calls will by default be automatically forwarded to.

Enable Voice Prompt. If enabled this feature causes a voice prompt to be played immediately prior to forwarding the call. The name of the voice prompt to be played (see Voice Prompt Management later in this section) can also be configured.

10 In addition to the above configuration parameters, the number to be dialed can be overridden by binding the parameter “Autoforward_CallNumber” to the number that is to be called. This allows a different number to be forwarded to based on DNIS/DID information, or other call setup information.

15 e. ***Internal Services***

 The gateway server, by default, has several “internal” services that are used to accommodate common tasks, such as collecting digits and obtaining user identification. These services are not directly selectable, but can be used when developing third party applications, and may also appear in the list of services that are presently in operation.

4. ***Call Rating***

 The gateway server is capable of assessing charges to user accounts for calls that are made by the user through use of the Call Rating Manager 358. These charges can either be recorded in call detail records, or debited directly from users accounts. This process is explained in greater detail when the gateway application, and User Management are explained. However, a key part of performing this charging for a call is determining how much to charge for the

call. This section explains the operation of call rating, and describes how call rating can be configured. The Call Rating Manager 358 of the gateway server is designed to be extensible by application developers. In order to facilitate this, there are multiple ways to rate calls, known as call rating types. A single standard model for rating calls is built into the software; it is possible for third party application developers to add additional means of rating calls. In addition, for each call rating type, numerous profiles may exist. Different profiles may use different values for how much to charge, how often the user is charged, and what levels of discounts are applied. The call rating engine supports multiple profiles for each of the call rating types that are present.

The call rating configuration dialog 9002 can be found under the “General” node in the configuration tree. Expanding this node will show the “Call Rating” node, which can be double clicked on in order to bring up a call rating configuration dialog. This dialog is shown in FIG. 90 and described below. The billing profiles panel 9004 provides a list of currently defined call rating profiles for each call rating type. New profiles can be added to the list, and existing entries can either be removed or modified.

When modifying a call rating, the dialog that will be displayed depends on the call rating type. For the built-in call rating type, the following information is included in a Standard panel 9102 shown in FIG. 91 and described below:

General panel 9104 contains:

Description. Allows a textual description, for reference purposes, to be associated with the call rating profile.

Rate Modifier. Routes that are used to place calls contain a rate modifier, which is a floating point number. How this number is used for a given rating profile depends on this configuration setting. The rate modifier can be ignored, in which case the rate modifier field in the route will have no effect.. If the base option is selected, then the rate modifier will be used as the base rate for the purposes of rating the call (base rates are explained below). If the percentage

setting is selected, it will be used to adjust the normal base rate to determine the effective base rate.

5 *No Charge Time.* This field specifies the time, in seconds, during which no charge will be applied for the call. Calls of this duration or less will effectively be free. This can be set to zero to begin charging immediately.

10 *Base Rate.* This field specifies the base rate for the call, in units (dollars) per minute. The base rate is the nominal charge that will be applied during the entire call. This field is overridden if "Base" is selected for the Rate Modifier, and will be adjusted by the percentage specified in the rate modifier of a route if "Percent" is selected for the rate modifier.

15 *Granularity.* This controls the incremental billing unit. It is used to ensure that all charges are in cents, tens of cents, dollars, etc. This allows currencies such as the Yen, which has no fractional components, to be used, by simply setting the granularity to 1 (i.e. all billing is in increments of 1 Yen).

20 Call Rating Setting panel 9106 contains:

25 *Initial Interval.* Specifies the first interval, in seconds, which will be charged for when making the call. As soon as the no charge time is exceeded, the initial interval will be billed for.

30 *Initial Rate (%).* This controls the actual rate that will be used for billing the initial interval. For example, if the initial rate is set to 200%, and the initial interval is 10 seconds, then the first 10 seconds of chargeable time will be billed at a rate of 200% of the base rate. This can be used to provide certain behavior. In the example above, if the no-charge interval is 10 seconds, the initial interval is 10 seconds, and the initial rate is 200%, then the first 10 seconds are free, but
35 after 20 seconds, the charge is for 20 seconds worth of the base rate. By doing this, calls are effectively charged at a uniform rate, but if the caller hangs up in the first 10 seconds, no charge will be applied.

40 *Incremental Interval.* Once the initial interval has elapsed, the user will be billed in incremental units of time specified by this parameter.

Incremental Rate (%). Each incremental user billed to the user will be billed at that incremental rate percentage times to the base rate for the call. Normally, this will be 100%. As an example, if the incremental interval is 10 seconds, the base rate is 0.60 per minute, and in the incremental rate % is 100%, then the user will be billed for each 10 seconds used, or part thereof, at 100% of the base rate of 0.60 per minute. Since each 10 seconds used costs 0.10 at 100% of the base rate, each 10 seconds or part thereof used by the user would cost them 0.10.

5. *Voice Prompt Management*

An integral part of providing IVR based services is the management of the voice prompts that are actually played to the user. The voice prompt management system is responsible for managing the voice prompts that are available on the system, and allowing those voice prompts to be played. In addition, in order to support multilanguage applications, multiple languages of a particular voice prompt may need to be supported.

The Media Manager 356 handles all aspects of voice prompt management. It allows the list of voice prompts to be configured via the management system. To view a Media Manager window 9202, double click on the "Media Manager" node in the configuration tree, which is visible when the "General" root level node is expanded. The Media Manager window is shown in FIG 92.

The Media Manager window 9202 shows the list of voice prompts in Media Manager Settings panel 9204. By default, when ExpressConfig sets up the default configuration for the gateway server, a number of voice prompts are installed. These voice prompts are used by the gateway application in order to interact with the user. A user can modify these voice prompts to change the prompts used by the gateway application. One can also add additional voice prompts that can be used by third party applications, or other applications such as Autoforward

Media Manager buttons for Adding and Removing function differently than a normal list. The reason is that each entry has a number of versions (i.e. languages) of each voice prompt may be present. When one adds a new media entry, New Media Entry dialog 9302 for adding a new voice prompt is shown in FIG. 93.

Settings panel 9304 allows for configuration of the prompt. The name of a voice prompt should be specified. If one selects the name of an existing voice prompt, one should add an available version of that voice prompt (or overwrite an existing one). If one enters a new name that does not already exist, a new voice prompt will be created and maintained by the media manager.

The version corresponds to the language selection of the user, and can also be used for other purposes. When adding a voice prompt or version of a voice prompt, one can also specify the default version of a voice prompt. The default version is the version that will play if the requested language is either unspecified, or unavailable. The description of the voice prompt is used for informational purposes. Finally, one needs to provide a path to a muLaw file that contains the voice data that one wishes to be played. This path is specified in Media Entry File panel 9306 in FIG. 93. This file should be in raw muLaw format, 8000 samples per second, 8 bits per sample, mono.

Similarly, when one removes a voice prompt, a user should select which version is to be removed. One will be provided with a selection of the available versions that can be removed. If one removes the last version of a particular voice prompt, then that voice prompt will be removed entirely from the system.

If one needs to make bulk changes to voice prompts, it is useful to be able to directly access the voice prompts as they are stored on disk. The filename format for voice prompts managed by the media manager is

`mm_<promptname>_<versionname>.mulaw`

where <promptname> is actually the name of the voice prompt, and <versionname> is actually the version of the voice prompt. The files are stored

in the main gateway server directory. Using this information, one can do a bulk overwrite of all voice prompts once one has a configured set of voice prompts. (However, this can only be used to overwrite existing voice prompts, not to add new voice prompts or new versions of voice prompts). One should also note that one should restart the gateway server if her or she directly modifies the files on disk in order for changes to take effect.

2. *The Gateway Application*

This section further discusses the gateway application described in the previous section. The gateway application is the main application for delivering calling services on the gateway server platform. It makes use of all of the capabilities provided by the gateway server platform in order to deliver sophisticated calling services to the end user.

The discussion of the gateway application is divided into several parts, as listed below.

Principles of Operation – Describes the basic sequence of steps that the gateway application uses in processing a call entered by the user. This allows the basic IVR sequences to be understood. It also allows an understanding of the steps that the gateway application goes through in setting up a call.

Gateway Application Configuration – Discusses the configuration options that can customize how the gateway application behaves, and how to set those parameters.

Supplementary Application Configuration – Discusses the internal applications that affect the behavior of the gateway application, and how these applications can be customized.

Session Bindings – Describes the session bindings interpreted by the gateway application.

Each of these parts is discussed in a separate section below.

1. Principles of Operation

The gateway application is a separate application from the gateway server that interfaces with the gateway server using the services API provided by the gateway server. However, the gateway application is normally loaded in-process by the gateway server. This allows the gateway application to integrate with the management and configuration facilities provided by the gateway server, in addition to being able to provide services to end users.

The gateway application handles calls, and interacts with users in order to collect the necessary information to complete calls on behalf of the user. The gateway application handles both network calls and telephony calls, and handles inbound or outbound calls of each of those types.

The following sequence of steps is followed by the gateway application in handling a telephone call from a user, as shown in FIG. 13:

- The incoming call is answered (step 1315);
- The system greeting is played (not shown, optional);
- Depending on the configuration, the user may be prompted for their identification (step 1320). If so, the identification entered by the user will be processed to ensure that it is correct, and the call will proceed (step 1330);
- The user is requested to enter the number that they wish to call (step 1332);
- The number dialed by the user is translated using a E.164 translator (also called a E.164 Parser) (described in the section on routing and number translation below) into a standardized format (FIG. 18);
- The Routing Manager 352 is used to find routes that match the number dialed by the user. The routes can come from a number of sources. The process used to find routes is discussed below;
- The routes are prioritized according to the configuration of the system, as described in the routing model section below, various types of routes can be implemented to accommodate user requirements;

– Routes are tried successively, until either all routes have been tried, or the call has been placed successfully. If the route being used requires the user to provide their identification and/or authentication code, then the user will be prompted to do so before the route is actually used,

5 – If debit is being used for the call, then a debit reservation is made against the user account that is making the call. Depending on the configuration of the gateway service, the user may also be informed of their current balance, and how long the call will be allowed to last,

10 – The call is placed according to information contained in the route. If a connection is established to the remote party, then the call proceeds as normal. Otherwise, the system plays a message indicating why the call could not be completed and optionally prompts the user for another number to call (for instance, step 1445 of FIG. 14);

15 – If the call is connected successfully, and the called party hangs up, the calling party will be informed, and optionally asked for a second number to call. If the calling party hangs up, the called party will be disconnected;

20 – After the call is completed, regardless of whether or not it was successful, a call detail record will be inserted into the database of call records. A cause code will indicate whether or not the call was successful (and why).

2. *Gateway Application Configuration*

20 Several configuration parameters can control how the gateway application interacts with the end user. These parameters are explained below. The Gateway Application Window 9402 can be viewed by opening the “GatewayService” node in the configuration tree, which is visible when the “Services” root level node of the configuration tree is expanded.

25 There are three panels accessed from Tabs 9404 that allow the configuration of the gateway application. These panels are shown in FIGs. 94 -

96. The tabs 9404 at the top of the window allow these panels to be navigated between.

a. General Settings Panel

General Settings panel 9406 allows general gateway application configuration options to be configured. There are only three parameters that can be configured on this panel:

Follow on calls. This parameter determines whether or not the user will be allowed to make follow on calls. This refers to the capability of placing a second (and third, and so forth) call after the first call has been disconnected, without hanging up and calling back into the system. If the gateway application serves live users, this is generally a useful feature; however, if the gateway server is interfaced to an automated switch, then this functionality is generally not necessary.

Maximum number of retries. To limit the length of time that a particular line can be tied up, there is a limit to the maximum number of retries that can be entered by a user. This field controls how many times a user can try to unsuccessfully make a call before they will be disconnected by the system.

Call Charge Preferences. This field allows the selection of the strategy for charging calls placed using the gateway application. The possible selections are as follows:

– Credit. Calls will be made in a post-billed mode. Debit will not be used for calls, and calls will not be made without a charge applied;

– Credit/Debit. Calls will be made in a post-billed mode. However, if the user of the system does not have credit enabled on their account, but have debit enabled, then debit will be used to place the call;

– Debit. Calls require that the users have sufficient debit balance in order to make the call. Users with insufficient debit balance, or whom do not have a debit account will be denied permission to make the call;

– Debit/Credit. Calls will be made in debit mode, as described above. However, if a user does not have debit enabled on their account, but has credit enabled on their account, they will be permitted to make the call using credit;

– No Billing. In this mode, credit/debit checking will not be performed, and the user will be allowed to make calls without charges being assessed. However, CDRs will still be collected, so a post-paid billing system based on CDRs would still be functional;

– Route Dependent. In this mode, the call charge preferences depend on the number dialed, and are contained in the route.

b. Country/Area Code Settings Panel

The country/area code settings panel 9502, shown in FIG. 95, allows one to provide information about his or her local calling area to the gateway service. The gateway service is provided by the gateway server. It includes, but is not limited to, handling calls, bridging communications systems, and providing access to other systems, such as databases, routing servers, etc. This is used to provide a dialing plan that emulates the dialing plan provided by the local telephone company in that area.

There are four parameters that one needs to configure. As shown in FIG. 95, these parameters are the international prefix, national prefix, country code, and area code. The way that these parameters operate is as follows. After the user enters the number, it goes through the following three steps prior to being interpreted by the system:

– If the number is prefixed by the international prefix (i.e. the sequence of digits normally used to dial an international number), then the international prefix is stripped off, and the number is interpreted;

– If the number is prefixed by the national prefix (i.e. the sequence of digits normally used to dial a domestic long distance number), then the national

prefix is stripped off, the country code is added to the front of the dialed digits, and the number is interpreted;

– If the number is not prefixed by either of the above, then the country code, followed by the area code, are added to the front of the dialed digits, and the resulting number is interpreted.

For example, if the international prefix was 011, the national prefix was 1, the country code was 81, and the area code was 3, then numbers would be processed as follows:

– 01117035551234 matches the international prefix and is processed as 17035551234;

– 1612345678 matches the national prefix and is processed as 81612345678;

– 123456789 matches neither, and is processed as 81312345678.

In the first example, nothing was added after the international prefix, since dialing the international prefix implies that the country code and area code have been dialed. In the second case, the number didn't include the local country code (81), since it was a domestic long distance call – so the country code was added prior to processing. In the final example, neither the international nor national prefixes were added; so it is assumed that the call is to the local area, and the country code plus area code were added to the front of the dialed number.

When the dialed number is sent for processing and route resolution, it is always required to be in a fully resolved format that contains the country code, area code, and local number. This way, routing information can be shared between gateway servers in different areas without modification. If routes were simply based on digits dialed by users, locale differences would prevent this from being possible.

c. ***Advanced Panel***

Advanced panel 9602 contains settings that do not normally require modification. The settings available on this panel are shown in FIG. 96 and discussed below.

5 *Try First Possible Route.* By default, the gateway application will try to place a call using the optimal route, and ask for user identification if necessary. If some users cannot be identified, this could cause difficulty since even if multiple routes exist, they will not be tried if the user cannot enter their identification code when the first route is tried and requires identification. To circumvent this situation, it is possible to enable the first possible route (that can be used with the current credentials) to be tried first, rather than the optimal route.

10 *First Digit Timeout (ms).* This parameter controls how long the gateway service will wait for the first digit to be dialed when waiting for a number.

15 *Subsequent Digit Timeout (ms).* This parameter controls how long the gateway service will wait for the second and subsequent digits when waiting for a number.

20 *Identification Mode.* Selects when identification information will be requested from the user. There are three possible settings for this parameter:

Always – Always prompts the user for their identification, after the number has been entered, but before the call is actually placed.

At the beginning – Always prompts the user for their identification, but does so immediately, when the user first calls into the system.

On demand – Only prompts the user for their identification if it is required in order for the user to have sufficient privilege to access.

25 ***3. Supplementary Application Configuration***

In order to facilitate reuse of certain application code by third party application developers, the gateway application separates some functionality into separate applications that are used to perform common tasks, such as obtaining

a valid user identification, or obtaining a users authentication code. These supplementary applications are separated from the gateway application, and are invoked by the gateway application. This section describes the configuration of supplementary applications invoked by the gateway application.

a. User Identification Application

The user identification application, part of Higher Level Services 315, is used to obtain a valid user identification from the user. The user will be prompted for their user identification, and will be given several attempts to enter their user identification correctly. The user identification application will automatically consult with the database server to determine if the user identification provided is valid. The user identification settings window 9702 can be accessed using the "IdentificationService" node in the configuration tree, visible once the "Services" node is expanded. Window 9702 is shown in FIG. 97 and displays Identification Service Settings panel 9704 which relates the following information:

Maximum Number of Retries. This parameter controls how many failed attempts to enter their user identification a user will be allowed before they are disconnected.

First Digit Timeout. This parameter determines the amount of time a user will be given to enter the first digit of their user identification.

Subsequent Digit Timeout. This parameter determines the amount of time a user will be given to enter the second and subsequent digits of their user identification.

b. User Authentication Application

The user authentication application, part of Higher Level Services 315, is almost identical to the user identification application, except that it obtains and checks the authentication code for a user. Configuration screen 9802, shown in

FIG. 98, can be accessed using the “Authentication Service” node in the configuration tree, also visible once the “Services” node is expanded. Authentication Service Settings panel 9804 is displayed and relates the following information:

5 *Maximum Number of Retries.* This parameter controls how many failed attempts to enter their user authentication code a user will be allowed before they are disconnected.

First Digit Timeout. This parameter determines the amount of time a user will be given to enter the first digit of their authentication code.

10 *Subsequent Digit Timeout.* This parameter determines the amount of time a user will be given to enter the second and subsequent digits of their authentication code.

4. **Session Bindings**

15 The previous section discussed the concept of session bindings, in particular how they relate to customizing how a service application behaves. This section discusses the session bindings that can be used to manipulate the behavior of the gateway application.

20 The gateway application uses several session bindings that are set up automatically by the gateway server platform in order to provide the appropriate behavior based on the type of call being handled. At this time, there is only one parameter that can be directly used to customize the behavior of the gateway application. This parameter is shown in the table below.

Parameter	Value
UserLanguage	Version of media prompts that will be played by default, until the user is identified. Once the user is identified, voice prompt language will be determined by the user profile

5. *Routing and Number Translation*

This section describes how to configure routing and number translation according to the present invention. This is a critical aspect of the overall operation of the VoIP system. There are numerous ways that calls can be routed across the network in the present invention, and selecting the optimal approach for a particular configuration is important in ensuring the performance and manageability of the network. This section presents the concepts behind routing and numbers translation, and describes how routing affects the functionality of the overall system, and explains available choices for how calls will be routed.

Routing in the context of the present invention software refers to call routing, not IP routing. IP routing, which handles the routing of IP packets sent on the underlying IP network, is a separate issue and it is assumed that the present invention components are running on a network in which the IP routing is configured correctly.

This section is divided into several sections, each covering a different aspect of call routing. These sections are as follows:

Number Translation. As a prelude to discussing actual routing, one should consider how the present invention software translates digits dialed by the user into a call address that can be routed. This section describes the number translation process.

Routes and Route Information. Describes the concepts of routes, and how routes are used in order to complete calls. The information that is contained in

each route, and how that information affects the behavior of the system, is also discussed.

Routing Models. Explains the different types of routing models that exist, based on how routing information is distributed in the network. Local routes, static global, disseminated, and dynamic routes are all discussed.

Routing Manager. Having discussed the principles behind a route, this section discusses how routes are managed, and general settings that are used to control the routing process.

Routing Servers. To facilitate the exchange of routing information between different gateway servers, the present invention provides routing servers. Routing servers can be used to centralize the administration of call routing in the network.

a. Number Translation

The gateway application normally collects digits from the user to determine what number they wish to call. However, the routing system is based on routing numbers in a specific format. Specifically, the routing system in the present invention is based on routing E.164 numbers, comprised of a country code, an area code, a local number, and an optional extension. These E.164 numbers, also called E.164 addresses, are usually displayed as follows:

+CC(AC)LN-EEE

In the above, CC would be replaced with the country code, AC with the area code, LN with the local number, and EEE with the extension. The routing system is capable of routing numbers in this format, so the string dialed by the user is translated into the format shown above. For example, the following translation would be used to make a call:

170378770007002 -> +1(703)7877000-7002

Although the above example is quite simple, certain countries have complex system for number translation. In these situations, the number of digits

in the area code and local number may be variable. The subsections below explain the method used by the system to perform number translation

b. E.164 Parser

The E.164 Parser (shown as element 14708 in FIG. 147) is part of Subsystem 351 of the Routing Manager 352 in the gateway server. The E.164 Parser is responsible for translating the digits entered by the user into E.164 form. The E.164 Parser operates using tables that are locally defined on each gateway server. However, it is possible to allow E.164 Parser data to be synchronized automatically, to prevent the need to update E.164 Parser tables on all gateway servers each time a change is to be made. The synchronization process is described in the next section and is shown in FIG. 18.

A parser subsystem is maintained locally on both the gateway server (as a subsystem of the Routing Manager) and the routing server. These two components maintain a connection with each other. Each parser maintains a listing of translation patterns (steps 1810 and 1815). Each parser maintains a local table for parsing addresses (steps 1820 and 1825). If synchronization is enabled (steps 1830 and 1845), the two components connect (steps 1840 and 1855) and the parser on the routing server sends new and/or updated parsing information/data to the parser on the gateway server (steps 1860 and 1870). In this way, numerous gateway servers can be synchronized with the latest parsing data from a centralized location.

The E.164 Parser operates by matching patterns in the E.164 Parser table with the digits dialed by the user. Each table entry specifies the pattern of dialed digits that should be matched, as well as information on how to translate those digits. Table entries are organized by country codes; table entries within each country should be set up for which numbers can be parsed. The gateway server is installed with a default set of tables, but E.164 parsing information should be provided for countries not set up by default. It is also necessary to update E.164

parsing information to reflect actual changes to the numbering system in the world.

FIG. 99 shows Country/Area code panel 9906 with an exemplary list of countries for which parsing data is configured can be viewed by expanding the "Routing" root level configuration tree node, and double clicking on "E.164 Translation". The synchronization mode, shown in panel 9904 and explained in the next section, can also be configured from this screen.

The country list is a standard list; selecting a country will allow its entries to be added or removed; country codes may also be added. Once a country is selected, a user is able to configure the table for that country, as well as configuration settings for that country. As shown in FIG. 100, general panel 10006 for a country allows one to enter the country code and country name for the country that her or she is configuring. General Panel 10006 is one of two Tabs 10004 in screen 10002. The country code cannot be changed, however, once the country has been added.

Address List panel 10102, shown in FIG. 101, allows the actual translation table for the country to be set up. All existing configured patterns are listed, and can be modified or removed. Additional entries can be also be added.

There are two pieces of information that need to be configured for each entry in the E.164 Parser table. The first is the actually pattern, specifying addresses that will match. This information is configurable under the Address panel 10206 when configuring a table entry, and is shown in FIG. 102. Address panel 10206 is shown in screen 10202. Screen 10202 has two Tabs 10204.

There are four pieces of information that are required when specifying an address to be matched:

Country Code. This is the country code to which the entry pertains, and is not editable since the country which entries are being added to is already selected. When checking a number entered by the user, the E.164 Parser starts at the left of the entered digits, and attempts to find a matching country code for

the first digit entered. If it fails, it tries again for the first two digits, then the first three digits, and so forth. Once a country code is selected, the E.164 Parser proceeds to the area code checking step.

5 *Area Code.* The digits following the country code will be compared against the area code field. If the digits following the country code match the value entered in the area code field, then checking will progress to the local number stage. Since it is not desirable to enter every conceivable area code, it is possible to wildcard the area code. Wildcarding is accomplished by entering 'A' in the place of a digit that is to be matched, and will allow any digit to be matched. For example, the pattern 4AA will match any three digit sequence starting with the digit 4. However, once one uses a wildcard digit, all digits to the right of that point should also be a wildcard. Using the above example, 4AA is a valid entry for the area code field, but 4A4 is not.

10 *Local Number.* If the area code field matches, then checking will proceed to the local number stage. At this point, the digits following the area code will be compared against the digits specified in the local number, and a match will be checked for. If the digits match, then the checking will proceed to the extension stage. As with the area code, it is possible to supply wildcard digits in the local number field, by using the symbol 'N'. Thus, 123NNNN will match any seven digit local number, as long as it begins with '123'. Additionally, it is possible to specify optional digits, using a lowercase 'n'. For example, 123NNnn would match numbers starting with 123, with between 5 and 7 total digits. Note that all characters to the right of an 'N' should either be another 'N' or an 'n', and that all characters to the right of an 'n' should also be an 'n'. Finally, if one uses optional wildcards, the extension field should be blank.

15 *Extension.* For certain applications, it is desirable to allow an extension to be explicitly represented. For example, a gateway server connected directly to a PBX whose main number is +1(703)555-1234 might want to terminate calls to 17035551234 followed by any four digits (assuming four digit extensions on

the PBX). The extension field is optional, and be left blank, in which case extension checking is bypassed. If present, the extension field specifies digits that must be matched in the digits dialed by the user that follow the local number. As with the local number field, it is possible to use 'E' as a required wildcard character, and 'e' as an optional wildcard character.

These wildcard characters form the basis of the call matching and prioritization features. FIG. 18 shows the E.164 Parsers on gateway server 210, 220 (although the parser, routing manager, and other components are available on all gateway servers). FIGs. 19 and 20 show the parsing and wildcard embodiment of parsing, respectively. In receiving call information data (step 1920), the E.164 Parser performs the steps of determining the country code (step 1930), retrieving the pattern information (step 1940), determining the area code (step 1950), determining the local number (step 1960), determining the extension (step 1970). During this time, the E.164 Parser can utilize the wildcard routine to assign priority and structure to the matching routes found above.

Wildcarding is shown in FIG. 20, and is described above with respect to the use of large and small letters to substitute for missing numbers in a E.164 address. Specifically, wildcarding allows for the prioritization of partially matched routes by sorting based on fewest wildcards of the lowest degree. This prioritization is described in detail in the next section. The parser does high level parsing of the route and provides the parsed address information (step 2005). The parser then attempts to match the values for the area code based on the parsed country code (step 2010). The parser then moves to the phone number (step 2020) and extension (step 2030) respectively. The results are prioritized based on the fewest wildcards (which correlates to the most matching digits)(step 2040). The results are outputted as matched E.164 addresses (step 1980).

The second panel under a pattern entry for a particular country, "Extra Settings", allows one to configure the number of extra digits that will be matched with the pattern. If a pattern matches the dialed number but additional contains

extraneous digits, then the extra digits setting will be consulted. If there are fewer extra digits than specified in this setting, then the pattern will be considered to match and the extra digits will be discarded. Otherwise, the digits entered by the user will not be matched.

c. Prioritization and Examples

Another important concept in understanding E.164 translation according to the present invention is that of prioritization. There may be multiple entries that match a particular number dialed by the user. The E.164 Parser selects the table entry that has the greatest number of matching digits. For example, if the number dialed is 17035551234, consider the following entries exist in the table for country code 1:

+1(703)55NNNN
+1(70)35NNNNNN
+1(703)NNNNNNN
+1(AAA)NNNNNNN

The matching table entry selected will be +1(703)55NNNN, because it has the most digits matching. Note that although +1(70)35NNNNNN also matches the dialed digits, fewer specific digits match (in other words, there are more wildcard substitutions). Also, note that the result of using +1(70)35NNNNNN would have been different, since it would have resulted in an area code of 70, and an eight digit local number.

FIG. 21A illustrates a routine for prioritizing routes. The E.164 Parser checks the route's address entry (step 2105), the route's timing information (step 2110), the route's access entry (step 2115) and the route's ordering information (step 2120). The E.164 Parser then determines which route(s) are available (step 2125) and compares a reduced list of routes (step 2130) with the address requested by the user. The parser then provides a prioritized list (step 2135).

5

CC	AC	LN	EXT	Description
1	AAA	NNNNNNN		Matches any 10 digit number in country code 1. The first three digits will be used as the area code, the following seven as the local number. Numbers matching this rule will have no extension.
1	703	555NNnn		Matches numbers starting with 1703555. Numbers may be 8 to 10 digits in length. The area code is always 703; the remaining digits dialed is considered the local number.
1	703	5551234	eee	Matches 17035551234, with up to three additional digits. If present, the additional digits will be considered to be the extension field.
81 81	429 42	5NNNN 91NNNN		8 digit (excluding country code) numbers beginning with 814295 will be considered to be in the 429 area code, whereas numbers beginning with 814291 will be considered to be in the 42 area code. This true example (taken from Japan) illustrates how local number digits can actually alter the determination of what area code a number falls into.

d. Alternative Configurations

The sections above discussed how E.164 pattern tables are used to translate numbers to an E.164 form. The E.164 Parser system is based on emulating the telephone number interpretation typically performed by PSTN switches. However, this is not necessarily appropriate for all applications. For example, when terminating traffic to a foreign area with variable length area

codes, it may be known only that all numbers are 9 digits in length when the area code and local number are combined. In such situations, it is possible to forgo specifying the area codes, and simply to treat the entire sequence of digits entered by the user as a long local number. In this case, the area code field would be left blank, and the local number would be 'NNNNNNNNNN' (9 wildcards). With this entry, it would not be possible to know the area code that was actually dialed, but this knowledge is not always required, and using this approach can avoid the otherwise tedious task of creating a list of all existing area codes.

Note, however, that any routes configured on the gateway servers would also have to be for a nine digit local number with no area code.

e. E.164 Parser Data Synchronization

The above sections described a key element, the E.164 pattern tables, that are used to control how digits entered by the user are parsed into a standard address form. However, the tables that are used in doing this are relatively large, and furthermore, may change frequently if they are used to reflect the way that numbers are translated by the PSTN, since PSTN carriers make frequent changes to the ways that area codes are assigned. In order to ensure that gateway servers continue to be practical to maintain, it is possible to synchronize E.164 Parser data with a centralized source. Whenever a change is made to the centralized version of the data, maintained on a routing server (discussed below), the change will automatically be propagated to the gateway servers that are synchronized with that version of the parser data, as shown in FIG. 18 (discussed above).

On a gateway server, E.164 Parser data synchronization is configured on the E.164 Translation node, described above in the section on the E.164 Parser. In the box labeled "Synchronization Mode", there is a button labeled "Change" that can be used to control how synchronization is performed. There are three options:

No Synchronization. In this mode, E.164 Parser data must be configured independently on the gateway server. Changes to E.164 Parser data elsewhere in the network will not affect the way that the gateway server being configured parses E.164 numbers.

5 *Client Synchronized.* In this mode, E.164 Parser data will be synchronized with a server that provides E.164 Parser data. This mode should only be selected on gateway servers, not routing servers (described below), except in specialized cases. In this mode, the contents of the local E.164 Parser tables will be kept consistent with those of a listed server. Whenever a change is made to the data on the server, the local gateway server will automatically receive those changes and update its local database. Because the local database is maintained, the number translation process remains very fast, so that overall call setup time is not impacted. The IP address and port number of the server are required when selecting this mode of operation.

10 *Server Synchronized.* This synchronization mode is not generally useful for gateway servers. For routing servers, described below, this mode allows gateway servers to synchronize with the contents of the local E.164 Parser tables. Any changes made to the local tables will be propagated to the gateway servers that list the routing server being configured as their server during Client Synchronization mode setup. Additional setup needs to be performed, in order to control which gateway servers are allowed to connect and retrieve data; this is described when the routing server is discussed.

6. *Routes and Route Information*

25 An important concept in understanding how calls are routed by the system is that of a route. A route, at a simple level, is a grouping of information that tells the gateway server how to make a particular call. For example, a route could tell the gateway server to make calls to a particular area code via the gateway server located at IP address 192.168.1.3. Similarly, a route could tell a gateway server

that in order to make calls to a particular number, it should use a particular set of telephone lines, and dial a specific sequence of digits. In general, most calls involve the use of two routes. On the originating Gateway, a route will exist supplying the address of the terminating Gateway. At the terminating Gateway, a route will exist supplying the actual number to send to the attached telephone equipment, and allowing the call to be placed. The former being a network route and the latter being a telephony route.

The information actually contained in a route is significantly more complex than the simple description above. Information on prioritization, timing, and access control are also contained in a single route. This section describes the key information contained in a route.

There are four main groups of information that are contained in every route that can be configured, which deal with different aspects of routing. These groups are as follows:

- The list of E.164 address patterns that the route corresponds to, and how the gateway server should place a call in the event of a match. For example, a route for +1(703)NNNNNNN should match the number +1(703)5551234 entered by the user, and may tell the gateway server to complete this call via a gateway server located at IP address 192.168.1.3;

- Ordering information that controls the order in which routes are tried, if multiple routes are available to complete a particular call. This information is necessary in providing redundancy. Information of this nature would allow two routes to a particular number to be defined. The less expensive route would likely be prioritized so that it will always be tried first, and only if it fails will the second, more expensive route be selected;

- Timing information, defining when a particular route may be used. For example, consider a technical support center with a branch on the east coast and a branch on the west coast of the United States. One route could direct calls to the technical support center between the hours of 9:00am and 3:00pm Eastern

Standard time to the center on the east coast. A second route would direct calls between 3:00pm Eastern Standard time to the center on the west coast;

– Access Information or Access Control Information, which is used to determine which users are allowed to make use of the route. As an example, calls to a customer signup center at +1(800)2SIGNUP might accessible to all users, but calls to all other numbers can require that the user have an active account. Access control information allows such conditions to be specified.

Additional information may be contained in a route, depending on the type of the route. The sections below discuss the information that is common to all routes, regardless of their type. The screens shown in FIGs. 103 - 108 are for locally defined routes. However, the panels are common to other types of routes as well. One may expand the “Routing” root level configuration tree node and double click on the “Local Routes” window. From the local routes window, one can either add a new route, or modify an existing one.

a. Addressing Information

The routing address panel 10306 shows the list of addresses that will be matched by a particular route. For example, if an address matches +1(703)NNNNNNN, then whenever the user dials a number with country code 1, area code 703, and 7 digits in the local number field, it will be handled by this route (unless there is another matching route of higher priority). FIG. 103 shows the address list panel in screen 10302. Tabs 10304 allow access to other panels discussed in detail below.

Under normal conditions, a route will only have a single address entry associated with it. The reason that additional address entries are allowed is to provide flexibility in addressing very specific applications.

Remember that E.164 translation happens prior to routing. Therefore, numbers will already be in +CC(AC)LN-EXT format before the routing process commences. This provides increased flexibility, since it is not necessary to

specify how the number is to be interpreted. For example, a route could exist for +1(Aaa)NNNnnnn. This would match any number in country code 1, with a 1-3 digit area code and a 3-7 digit local number. It is not practical to do this at the E.164 translation stage, since a number like 1703555 could be interpreted as +1(703)555, +1(70)3555, +1(7)03555, and so forth. However, at the routing stage, it makes addressing highly flexible; even a country with a complex area code system like Japan could be routed to with a single route for +81(Aaaaa)NNNnnnnn.

When a new route address entry is added, or an existing route address entry is modified, Routing Entry Properties dialog 10402 will appear for configuring the route address entry. The fields in this dialog are shown in FIG. 104 and explained below:

Panel 10404 contains:

Call Address. This group of fields specify the address pattern that must match the address entered by the user in order for this route to be selected. In one example, only E.164 addresses are supported; however, other types may be supported. For E.164 addresses, it is necessary to provide the country code, area code, local number, and extension of the addresses that should be matched. The same wildcards described when E.164 translation was discussed could be used. However, as discussed above, some restrictions regarding the use of wildcards are relaxed. Specifically, using optional wildcard characters does not require that all subsequent fields in the address use only optional wildcards. For example, 'Aaa' can be used in the area code field without requiring that the local number field consist solely of 'n's. However, restrictions within a field (fixed digits must precede required wildcards, which must precede optional wildcards) still exist.

Panel 10406 contains:

Line Group Type/Instance. When the gateway server is placing a call using a particular route, it is necessary for a line group to be selected in order to make that outbound call. The line group type and instance allow a particular line

group to be identified, in which case calls using this route will only be placed using that line group. In addition, if the line group type is set to dynamic, then an appropriate line group will automatically be selected based on the contents of the route, and availability of individual lines.

5 *Rate Modifier.* The gateway server is capable of rating calls and charging for those calls. The section on call rating in the gateway server Capabilities section describes the rate modifier, and how it is used in determining the cost of a call to the user. The value entered in this field as the rate modifier will be passed to the call rating function when determining the charge for a call.

10 Panel 10410 contains:

15 *Disable Route.* In order to allow calls to specific numbers to be restricted, it is convenient to create routes that prohibit access to the number instead of allowing it. The 'Disable Route' field is used for this purpose. A disabled route will be treated like a normal route, and prioritized as if it was a normal route. However, when the list of routes is prepared, all routes lower in priority than the highest priority disabled route will automatically be rejected. For example, it is possible to reject calls to all numbers beginning with 911 in North America with just a single route for +1(AAA)911nnnn with Disable Route selected.

20 *Device Type.* Two kinds of underlying devices can be used to place a call based on a route; telephony devices, and network devices. The kinds of information used by a network device to place a call is obviously different than the information required by a telephony device. This field allows the type of routing information being supplied to be specified. Note that specific fields may change depending on the type of device selected. This field could sometimes be referred to as the route 'type' – e.g. network routes and telephony routes – but is not referred to this way herein to avoid confusion with route distribution types (local, dynamic, etc.).

25 *Remapped Number/Station Address.* This field appears for Telephony Routing device types only. The remapped number field allows one to specify the

digits that will actually be sent to the connected telephone equipment when a call is placed using this route. The remapped number may contain the same wildcards used in specifying the address of the call, but may not exceed the number of wildcards contained in the address. For example, if the area code field reads 'AAA', then a remapped number of 'AA' would result in the first two digits (leftmost) of the area code being sent to the connected equipment. A remapped number of 'AAA' would send all three area code digits to the connected equipment, and a remapped number of 'AAAA' would be invalid. Placing optional wildcards in the remapped number will result in a substitution if a matching digit exists in the actual address, and no digits if no such digit exists. The format of the station address field is identical. However, the station address is used to select a particular line for placing the call, and does not affect the digits sent to the connected equipment.

Panel 10408 contains:

Host IP/Port. This field appears for Network Routing device types only. This field is used to specify the IP address and port number of a gateway server or H.323 terminal that will complete calls to a particular number. The default port numbers should generally be used unless otherwise required.

Panels 10412 and 10414 contain:

Enable Voice/Fax. Controls the kinds of calls that will be enabled. These fields are optional. Voice calls should always be enabled; voice calls are dynamically switched to fax when the appropriate tones are detected.

Proprietary/H.323. This field appears for Network Routing device types only. Since the gateway server supports multiple protocols for actually placing a network call, it is necessary to specify which protocol will be used. Calls handled by other present invention gateway servers should select the Proprietary protocol (or non-proprietary protocol), whereas calls to H.323 terminals such as Microsoft NetMeeting should select the H.323 protocol.

b. Ordering Information

The ordering information panel 10502 for a route, shown in FIG. 105, is very simple, and provides parameters that govern the order in which a route is selected. Two parameters can be modified through this page; preferability, and a cost estimate. Both of these parameters affect the order in which routes are selected; however, the exact relationship between these parameters and order is determined based on configurable parameters that are discussed under the General Settings section below.

c. Timing Information

It is possible to restrict the times of the day, as well as the days of the week, during which a particular routing entry may be used. Possible applications of this capability were already discussed above. FIG. 106 shows timing information panel 10602.

Route timing restrictions apply to use of the route, which occurs when the call is initiated. In one example, if a call extends beyond the time in the call route, it will not be automatically disconnected. For example, for a route enabled from 9:00am to 5:00pm, a call placed using that route at 4:59pm will be allowed to continue indefinitely past 5:00pm.

Use Timing. This flag controls whether or not route timing restrictions will be used. If disabled, a route can be used at any time of day. Otherwise, the time restriction table will be used to determine if the route can be used.

Use Local Time/Time Zone. If selected, use local time means that timing information in the route will be assumed to be for the local time zone. If such routes are transferred between different systems, then the timing information will be interpreted according to the local time on the gateway server that uses the route. This may or may not be the desired behavior. For example, consider a route is valid between 9:00am and 5:00pm, local time, distributed to two gateway servers, one on the east coast (Eastern Standard Time) and one on the west coast

(Pacific Standard Time). The east coast gateway server would allow use of this route between 9:00am and 5:00pm EST, and the west coast gateway server would allow use of this route between 9:00am and 5:00pm PST. However, if local time zone is not checked, then the time zone (relative to GMT) can be selected. For example, if the time zone in the example above was GMT-5:00 (EST), then the east coast gateway server would allow calls between 9:00am and 5:00pm EST but the west coast gateway server would only allow calls to be made between 6:00am and 2:00pm PST.

Allowed Usage Times. A list of allowed usage times is managed for each route that has timing information enabled. Each entry in the list specifies a time range during which the route may be used. An arbitrary number of entries may be added. Each range entry specifies a starting day and time, and a finishing day and time. If at least one matching timing entry exists at the time and day of week that a route is considered, then the route will be used.

d. Access Control Information

The access control panel 10702, shown in FIG. 107, allows restrictions on users that can and cannot use a route to be configured. The access control system is highly flexible, allowing a wide range of conditions to be implemented.

The list allows an arbitrary number of access control entries to be configured. When considering whether or not a particular user is allowed to make use of a route, the system begins with the first route in the list, and proceeds through the access control list until it finds a matching entry.

If no entries exist in the access control list, then access will be granted. Otherwise, if entries exist but no matching entries are found, then access will be denied.

For each route access entry, FIG. 108 shows some of the parameters that may be configured.

Route Access Type panel 10804 contains:

1848.0060001 - Array Telecom Corporation Confidential

Route Access Type. Each route access entry can either grant access, or deny access to user. Since the first matching entry in the route access list will be used to determine whether the route will be accepted or denied, it is possible to set up complex restrictions. For example, it would be possible to grant access to a particular group of users, but to deny access to one particular member in that group using two route access entries. The first, which occurs higher in the list, would deny access to the particular member. The second would grant access to the group. Since the first entry in the list is processed first, the deny access result would occur for the particular member for whom it was desired to deny access.

Call Charging panel 10806 contains:

Call Charging. This option determines how call charging should be performed for calls that are placed using this route, and when this route access entry was selected. For example, consider a user placed a call using a route, and granted access under an access entry that specified that Debit is required. In this case, the user would be required to have debit enabled on their account, as well as a sufficient balance, in order for the call to occur. A detailed description of the options for call charging are described in the section on gateway server Capabilities, in the section on Call Rating.

Call Rating panel 10808 contains:

Call Rating Type/Profile. This selects the call rating type and profile used to actually bill the call, if the call is to be billed. These parameters are also explained fully in the section on gateway server Capabilities, in the section on Call Rating.

Requirements panel 10810 contains:

Requirements. This list specifies conditions that should be met in order for a route access entry to be selected. If any of the conditions in this list are not true, then the route access entry will not be selected, and the next route access entry in the route access entry list will be considered. This is a standard list,

allowing individual requirements to be added and removed. There are several kinds of requirements that can be configured:

– GroupID; requires a specific value for the Group ID of the current user.

Can be used to grant access to a particular group of users.

– IsIdentified; simply requires that the user have an account on the system.

– UserID; requires a specific value for the User ID of the current user.

Can be used to grant access only to a particular user, such as a system administrator account for testing purposes.

– UserMinPrivilege; requires that the user have a certain minimum privilege level. This does not require that the user is identified, since all line groups have a default privilege level that is assigned to users that have not yet been identified.

7. Routing Models

From the previous section, it is clear that there is a significant amount of information contained in each route. Because of this, it is highly desirable to minimize the number of times that the same routing information must be entered. In order to provide flexibility while keeping administration to a minimum, the present invention software supports four models for how routing information can be distributed around the network. These models are not exclusive; in other words, one does not need to select a single model. Routes of all four types can be used on a particular gateway server.

a. Locally Defined Routes

Locally defined routes, more commonly referred to as local routes, are routes that are defined on a single gateway server. Shown in FIG. 14, they are not by default transferred to other gateway servers, or shared with a routing server. The route exists locally on that gateway server.

The benefit of a locally defined route is that it is simple to configure, since it does not require a routing server to be used. It is also efficient, since it is available locally to the gateway server, so it can be found immediately without making a request to a distant routing server. The Routing Manager 352 queries for matching local routes (step 1410), and is either provided with a route and provides it so that the call can be connected (steps 1415 and 1420) or it looks for routes elsewhere (step 1430). The drawback to a local route is that if there are a large number of gateway servers, it becomes difficult to administer local routes on each gateway server.

Steps 1430 - 1445 in FIG. 14 are replicated as steps 1550-1565 in FIG. 15, steps 1655 to 1670 in FIG. 16, and steps 1780-1795 in FIG. 17. These steps allow for various routing models to be used in combination. If the system does not find a matching route, it will attempt to use other routing models (step 1430). If other models are not available (steps 1440-1445), the system will provide a lack of service notice to the gateway which can pass this information along to other parts of the system. If other models are available, then they will be implemented (step 1435).

However, local routes are an important part of the dynamic route scheme, described below. Local routes can be exported to a routing server. Exported routes are not an exact duplicate of the locally defined route, and are considered to be a different kind of route. However, dynamic routes (exported local routes) are derived from the information contained in a local route. Dynamic routing is defined in more detail below, and how a local route is transformed into a dynamic route is described in the section on routing servers. Export configuration panel 10902 is specific to local routes. Shown in FIG. 109, panel 10902 determines whether or not a particular route will be exported and transformed into a dynamic route. This panel appears when creating or modifying local routes and includes the following information:

Enable Export. If this flag is selected, then the route will be exported to the routing servers listed in the export list.

Export List. This list contains the names of the routing servers to which the route should be exported. Routing servers and exporting routes is described further in the section on routing servers below.

b. Static Global Routes

Static global routes are routes that are statically defined on a routing server. Shown in FIG. 15, they are very similar to local routes, but since they are defined on a routing server, numerous gateway servers can query them. Static global routes are not downloaded to the gateway server; rather, when a gateway server wishes to place a call to a particular number, Routing Manager 352 on the gateway server will ask the routing server for any routes that match the number entered by the user (step 1510). The routing server will determine if there are any matches (step 1515) and return any static global routes that match the query (in addition to any dynamic routes, described below) (step 1525). The Routing Manager will then check for matches (step 1535) and provide them to the gateway server for connecting (step 1540). The Routing Manager will perform any other searches for which it is configured (steps 1550 and 1555).

The advantage of using static global routes is that it simplifies administration, even across a relatively large number of gateway servers. Static global routes generally require consistency between the configuration of routing servers and the configuration of gateway servers. For example, if a routing server directs all calls to a particular number to a gateway server (a network route), that gateway server will need a matching telephony route enabling it to place the call on the telephony network. In some cases, however, gateway servers should query the routing server at the time of the call in order to obtain those routes which can increase call setup time.

Setting up static global routes is described further in the section on the routing server below.

c. Disseminated Routes

Shown in FIGs. 9 and 16, disseminated routes are routes that are defined on the routing server, but distributed to gateway servers at the time that the gateway server connects.

FIG. 9 shows a flowchart diagram of the routine for disseminated routes. The routing server receives a request for route(s) (step 910). The routing server accesses the stored route(s) which are configured for dissemination to that Gateway (step 915). The routing server looks up route(s) specific to the requesting gateway server (step 920). The routing server then provides the route(s) to the requesting gateway server (step 930).

FIG. 16 provides a detailed flowchart of the steps taken by the various components of the system. First, the gateway server connects (step 1610) to the routing server. The routing server looks up the routes to be provided to this specific gateway server (step 1615) and provides the disseminated routes (step 1625). Similar processes are being performed on other Gateway and routing servers (steps 1620 and 1630). Unlike static routes, the gateway server does not query the routing server for disseminated routes at the time of the call; rather, it simply consults (step 1635). The matching process is roughly identical in that the Routing Manager searching for matching routes (step 1640) and provides information (steps 1645 and 1665) depending on the results of the search. The Routing Manager will perform any other searches for which it is configured (steps 1655 and 1660).

FIG. 10 shows a flowchart diagram of the routine for storing routes on the routing server for distribution to gateway servers. Routes are stored in tables which delineate the address entry (step 1010), the timing information (step 1015), the access entry (step 1020), the ordering information (step 1025), and the route's

5

10

19

Figure 1

20

25

exportation so that no dead or otherwise duplicative and/or out-of-date routes are retained and possibly used by the system.

The transformation process (step 650) is shown in block diagram form in FIG. 8. After receiving the local exported routes (step 720, 810), the routing server checks the route's address entry (step 820). The route's timing information (step 825) and access entry (step 830), as well as its ordering information (step 840) are also checked. The routing server then adds a unique identity to the route (step 850) and adds the identity of the exporting gateway server (step 860). The routing server also add a date and time stamp which denotes when the route was obtained from the gateway server (step 870).

A more detailed illustration of the present process is shown in FIG. 17. The Routing Manager 352 connects to pre-configured routing servers (step 1710). Any exportable local routes are exported (steps 1720 and 1725). The exported local routes are transformed into dynamic routes (steps 1730 and 1735). The Routing Manager queries for matching routes (step 1740) and is provided with matching dynamic routes (steps 1760 and 1765), which are cached by the Routing Manager (step 1755) in a dynamic cache discussed below.

There are several advantages to using dynamic routing. Firstly, it simplifies administration since routes can be shared between gateway servers effortlessly. Secondly, it does not have the consistency problem of static and disseminated routes whereby the pair of network and telephony routes must match. This is because the exported route will normally be transformed from a telephony route into a network route; but all other information in the route will remain the same. Thirdly, it improves the resilience of the routing process to the failure of a gateway server; if a gateway server fails, its routes will automatically no longer be available, and other gateway servers will route calls as if the gateway server that crashed did not exist. Normally, gateway servers would try to route to the failed gateway server, and only after a timeout occurred would they fallback to alternate routes. The disadvantage, however, of using dynamic

routes, in some cases, is that the gateway server must query the routing server at the time of the call, increasing call setup time.

It is not necessary to actually set up dynamic routes, since gateway servers export them based on information in a local route. The parameters that affect the transformation process are shown in FIG. 8 and discussed below, as well as in the section on routing servers which follows.

FIG. 8 shows step 650 from FIG. 6B. The Routing Server 230 receives exported local routes from the Routing Manager on the gateway server (step 810). The routing server checks the route's address entry (step 820), the route's timing information (step 825), the route's access entry (step 830), and the route's ordering information (step 840). The routing server then adds a unique identity to the route (step 850), the identity of the exporting gateway server (step 860), and the date and time that the route was obtained/exported (step 870).

e. Dynamic Caching

Above, the four types of routes were presented. Among those types, static routes and dynamic routes had the disadvantage that it is necessary to query the routing server on a per-call basis to obtain the routes for that call. This increases the call setup time, which may be critical in some applications. To reduce the impact on call setup time introduced by these two types of routes, the Routing Manager built into each gateway server is capable of caching routes that are obtained dynamically for the routing server. This feature is discussed when describing the routing manager, below.

f. Summary

The following table presents a comparison between the four different kinds of routes, and is a summary of their benefits and disadvantages in certain examples:

Item	Local	Static	Disseminated	Dynamic
Configured On	Gateway Server	Routing Server	Routing Server	Routing Server
Call Setup Overhead	No	Yes	No	Yes
Route is Shared With	No Sharing	All Gateways	Selected Gateways	All Gateways
Configure Pair of Routes	N/A	Yes	Yes	No
Auto-Disable on Failure	No	No	No	Yes

8. Routing Manager

The Routing Manager is a built-in component of the gateway server. It is responsible for making all routing decisions, and for managing all the routes used by the gateway server. The Routing Manager has the following responsibilities, as shown in FIG. 11:

- Maintaining the database of local routes (step 1110);
- Establishing and managing connections to the routing servers in the network (step 1120);
- Exporting local routes to the routing servers, and importing disseminated routes from the routing servers (steps 1130 and 1140);

- Obtaining static global and dynamic routes from the routing servers, and caching those routes for future use (steps 1150 and 1160) ;

- Finding all matching routes for a particular number dialed by the user, and prioritizing those routes (steps 1145, 1147, 1170 and 1180).

5 The main configuration screen for the Routing Manager is found under the “Routing” root level node of the configuration tree. Expanding this node will show the other options available. Window 11002, as shown in FIG. 110, appears as described below:

General Settings panel 11004 contains:

10 Clicking on the “Change” button allows the general settings to be modified. There are three parameters that can be configured, as follows:

Dynamic Routing Usage. Selects when a query is issued to the routing server in order to obtain routes dynamically. (This includes obtaining statically configured global routes, not just dynamic routes). There are five options, as follows:

15 – *Always.* If selected, routing servers (if any) will be queried on each call to obtain dynamic and static global routes.

 – *Never.* routing servers are never queried for dynamic and static global routes.

20 – *Use if no Cached Routes Available.* The routing server will be queried only if no routes in the dynamic cache match the number that was dialed by the user.

 – *Use if no Local/Disseminated Routes Available.* The routing server will be queried only if there are no matching local routes, and no matching disseminated routes.

25 – *Use if no Routes Available.* The routing server will only be queried if no routes exist that can be used to place the call. This requires that there are no matching local or disseminated routes, and the dynamic cache has no routes that match the number entered by the user.

– *Cache Timeout*. Entries in the cache need to have a maximum lifetime after which they are no longer valid, since cached routes do not reflect any possible changes that occurred on the server. This parameter allows one to specify the lifetime of routes stored in the dynamic cache.

5 – *Base Error Filename*. This parameter allows one to configure the base filename for error logging. Date and session information will be appended to form the final filename.

Route Ordering panel 11006 contains:

10 – *Route Ordering*. This setting allows control over how routes are prioritized. If multiple routes are available for a number entered by the user, these parameters will be used to control the order in which those routes are tried. Clicking “Change” allows the list to be modified. It is possible to move each parameter up or down to select the order in which priority will be determined. For example, if “By address” is the top entry, followed by “Cost estimate”, then
15 when comparing two routes, the route with a better address score (described below) will be ordered first. If, however, the two routes have the same address score, they will be ordered based on their cost estimates. If they are the same on the cost estimate, then the next parameter in the list will be used.

20 In ordering routes, “By Address” refers to how many exact digit matches occur in the address. For example, if the user dialed +1(703)5551234, a route to +1(703)555NNNN would receive a score of 7, whereas a route to +1(703)NNNNNNNN would only receive a score of 4. Higher scores are more preferable.

Routing servers panel 11008 contains:

25 *Routing Server List*. This list shows the routing servers that the Routing Manager interacts with, and shows the current state of the connection to that routing server. Routing servers can be added and removed, or their settings can be modified. When adding or modifying a routing server entry, a significant

amount of configurable parameters can be specified. These parameters are discussed below.

a. Local Route Configuration

Local routes were already described at length previous sections, and how to configure a local route has been explained. In order to access the local route list, double click on the "Local Routes" node, under the "Routing" root level configuration tree node. This will display the local routes screen 11102. Screen 11102 contains a list from which local routes can be added, removed, or modified. The local route list screen is shown in FIG. 111.

b. Dynamic Cached Routes

The Dynamic Cached Routes window (not shown), which can be viewed by double clicking on the "Dynamic Cached Routes" node under the "Routing" node, shows the dynamic routes that are currently cached by the Routing Manager. This screen provides caching information.

c. Routing Server Settings

When adding or modifying a routing server, a large number of configuration parameters are available. These configuration parameters are split into seven panels shown as Tabs 11204 in screen 11202 of FIG. 112, each of which are described below. Note when viewing the configuration for a routing server that it is necessary to click on the change button in order to modify the configuration.

(1) General Panel

The general panel 11206, shown in FIG. 112, is used to specify general information about the routing server. Each of the parameters is discussed below.

Server Identification panel 11208 contains:

1848.0060001 - Array Telecom Corporation Confidential

Server Name. Defines a local name for the server. This name is used when identifying which servers local routes should be exported to. Note that the server name cannot be modified once a server has been added.

IP Address/Port Number. Defines the IP address of the routing server.
5 The port number should be left at the default unless one is specifically running a routing server on a different port number.

Exporting/Importing panel 11210 contains:

Send Exported Routes to Server. This parameter can be used to prevent any routes from being exported to the server. If this parameter is not enabled,
10 then no routes will be exported even if individual routes indicate that the route should be exported.

Accept Disseminated Routes from Server. If checked, then any routes disseminated by this routing server will be accepted. Otherwise, routes will not be accepted from the routing server. If dealing with a routing server that is not
15 trusted, disseminated routes should not be accepted since they can instruct a gateway server to place long distance calls via the telephone equipment to anywhere.

Use Imported Route Parameters for Disseminated Routes. If checked, then the parameters from imported routes will be used when the routes are disseminated. Route parameters include ordering, access and other information.
20 Local routing patterns may differ from the place of the route(s) origin. If so, then this should be unchecked.

Search Order. Multiple routing servers can be configured. Search order controls the order in which routing servers are used to obtain routes dynamically.
25 All routing servers marked "First" will be searched first. If all routing servers fail to find any routes, then all routing servers marked "Second" will be searched. This progresses until either routes are found, or no more routing servers are available to be searched.

(2) Exported Routes Usage

The exported routes usage panel 11302, shown in FIG. 113, allows one to override the settings of local routes that are exported to the routing server. This is part of the transformation process that a local route undergoes when it is transformed into a dynamic route. Exported routes will usually have their usage information changed in Usage Info panel 11304. IP Address and port number information, as well as protocol preference can be set in the Network Route Setting panel 11306. The Routing Manager makes these changes before the route is sent to the routing server.

(3) Exported Routes Access

The exported route access panel 11402, shown in FIG. 114, allows the route access settings in a local route to be replaced. Normally, it is desirable to remove route access settings sent to the routing server, since each gateway server will determine what route access conditions should apply. Configuration of the route access settings is identical to the configuration performed when defining the local route.

(4) Exported Routes Ordering

When exporting a route to the routing server, it may be desirable to customize ordering information prior to actually exporting the route. The exported routes ordering panel 11502, shown in FIG. 115, allows one to do this and includes the following information:

Usage. Allows one to specify whether the ordering information in a route will be overridden.

Preferability. If one is overriding the ordering information contained in the route, this field allows one to specify the preferability parameter that will be applied to the exported route.

Cost Estimate. If overriding the ordering information contained in the route, this field allows one to specify the cost estimate that will be applied to the exported route.

(5) Imported Routes Usage

5 The import route usage panel 11602, shown in FIG. 116, allows filtering of the types of routes that will be accepted from a routing server. This can be used, for example, to ensure that telephony routes are not accepted from an untrustworthy routing server, since such routes could cause the gateway server that receives them to make long distance calls on attached telephone devices. FIG. 116 shows this panel. Panel contains the following:

Routes Info panel 11604 contains:

Routes Info — These fields allow you to specify which types of routes will be accepted from the routing server.

Usage Info panel 11606 contains:

15 *Usage Info* — This field allows you to configure how usage information in the route will be modified, if at all. Modification may be enabled to change the rate modifier field of the route. The other route usage information parameters cannot be specified for imported routes.

(6) Imported Routes Access

20 The imported route access panel 11702, shown in FIG. 117, is used to replace the route access information on routes imported from the routing server. Often, each individual gateway server may define its own users, and its own set of permissions on what calls can be made by the user. However, routes from the routing server will typically not contain access settings that implement the desired behavior. This panel allows route access information to be overridden to implement the local policy for route access. Configuring the route access information is identical to configuring the route access table for a local route.

(7) Imported Routes Ordering

The imported routes ordering panel 11802 is identical to the exported routes ordering panel 11502 and is shown in FIG. 118. However, it is used to modify imported routes obtained through a dynamic query.

9. Routing Servers

The above sections described the principles of routing in the present invention, and discussed the Routing Manager that is part of the gateway server. An important component that was referenced significantly in the above discussion is the routing server. The routing server is a critical component in facilitating the exchange of routing information between different gateway servers, and for centralizing and simplifying the administration of routes.

This section describes the routing server, and how it can be used to deliver the routing capabilities described above. The sections below describe different aspects of the routing server.

The main routing server configuration window 11902 can be viewed by double clicking on the "Routing" root node of the configuration tree, when connected to a routing server. Window 11902 is shown in FIG. 119 and allows client setup to be performed.

a. Routing Clients

The top part of window 11902 in FIG. 119 contains List of Clients panel 11904. Panel 11904 can be used to configure the list of routing clients that are allowed to access the routing server. Groups of Aliases panel 11906 contains a listing of groups of clients. This feature allows for easier management of clients. These clients are gateway servers that are set up to connect to the routing server. Unless a client entry exists for a connecting gateway server, the connection will be denied.

When adding or modifying a client entry in the client list, several settings can be configured. These settings are described below. The configuration of each client is divided into four panels.

(1) General Panel

5 General panel 12006, shown in FIG. 120, is used to specify general settings regarding the client. General Settings panel 12008 contains:

10 *Disseminate Routes to this Client.* This field controls whether or not any routes can be disseminated to this client, and must be enabled to allow routes to be disseminated. If disabled, this parameter overrides the dissemination settings in individual routes.

Client Name. Provides a descriptive name for the client. This name is for information purposes only.

15 *Local Alias.* Routing clients may have a local alias. Multiple clients may share a single local alias. The purpose of this local alias is to provide a simplified means for identifying a client, or multiple clients. Aliases are the first level of grouping, and all clients are required to have an alias. If one does not require the first level of grouping, then one should simply use the client name as the alias name.

20 *Description.* The description allows text to be entered for informational purposes only. This information is not used by the system.

 List of Identities used by this client panel 12010 contains:

25 *List of Identities.* This field contains the list of identities that are associated with this client entry. The configuration of the identities is identical to the configuration explained herein when setting up the management client access list. The same procedure should be followed when setting up a routing client entry; the only difference is that this list controls the gateway servers that can access the routing client for routing purposes. Placing an identity in this list does not allow that entity to connect for administrative/management purposes.

Groups Belonging To. This list provides an easy way to add the client being configured into routing groups, which are described below. It also provides a simplified way of determining which groups the client is already a member of. However, it is necessary to take care when modifying group membership, since ALLASES are members of groups, not individual clients; thus, removing group membership removes the membership of the entire alias.

(2) Routing Usage Panel

The routing usage panel allows usage information in routes exported by the client to be replaced with specific routing information. Normally, this is used if the exporting gateway server is not fully trusted, and it is desired to ensure that all exported routes received from that gateway server are directed to the IP address of the exporting gateway server. The operation of this panel, however, is identical to the one found when configuring the Routing Manager component of the gateway server, and is described in the previous section.

(3) Routing Access Panel

The routing access information supplied in routes exported by individual gateway servers may not match the routing access information that it is desirable to export network wide. For example, a gateway server might define route access settings that only apply to its local users; the network wide scheme might specify different route access settings. Routing Access panel 12106, an example of which is shown in FIG. 121, allows the route access information supplied in routes exported by the gateway server to be overridden using settings that are defined on the server. The route access list is defined exactly the same way as for a local route.

(4) Exported Routes Ordering

Similar to usage information, it is possible to override the ordering information supplied in a route exported by a gateway server. This panel allows this to be performed in an identical manner to the way it is performed by the Routing Manager on the gateway server side.

b. Routing Groups

The routing server allows the definition of groups of aliases. Wherever an alias name can be supplied, such as in the list of gateway servers to which a route can be disseminated, it is possible to use a routing group name instead. This permits large scale operation of the routing server. For example, consider a routing server that manages routing across four carriers each of which has 10 Gateways in each of three countries. This is a total of 120 gateway servers. Each one would have a client entry, allowing it to connect to the routing server. For each carrier, three aliases would be defined – one for each country served by that carrier. At the group level, each carrier would have the aliases for each country grouped into a single group entry for that carrier. This would allow a route exported by one carrier to be available by the three other carriers, but would still allow a particular route to be disseminated to all the Gateways in one particular country.

Routing groups can be added and removed from the general panel shown above. The only thing required to create a routing group is the name of the group. Once created, members can be added or removed from the routing group

c. Exported Routes

The Exported Routes window 12202 is available by double clicking the “Exported Routes” node under the “Routing” root node in the configuration tree. This window displays all exported routes that are currently present on the routing server, and the name of the gateway servers that exported those routes. Window

12202 is shown in FIG. 122 and described below. One should not make any configuration changes using this window, but it is possible to use this window to verify that routes are being exported successfully.

d. Static Routes

5 Static Routes can be configured on the routing server by double clicking the “Static Routes” node, located below the “Routing” root node in the configuration tree. As shown in window 12302 in FIG. 123, the list of routes statically defined on the routing server will be displayed. The route list is a standard list that allows routes to be added, removed, or modified.

10 When adding or modifying entries in the route list, the route properties page is identical to that for a local route, except that the export list is not present.

e. Disseminated Routes

15 Disseminated routes can be configured in the same way as Static Routes. The dynamic route list is available by double clicking on the “Disseminated Routes” node, located below the “Routing” root node in the configuration tree. The Disseminated Routes window 12402, which is almost identical to that for static routes, is shown in FIG. 124.

20 When adding or modifying routes that are contained in the disseminated routes list, additional Properties window 12502 will appear that allows control over the clients to which the route will be disseminated. Panel 12506 is shown in FIG. 125 and includes the following information:

Disseminate to All. If selected, this flag will cause routing data to be disseminated to all clients that do not explicitly disable dissemination of routes.

25 *Dissemination List.* This list provides the name of all aliases and groups that are defined on the routing server, to which this route will be disseminated. Unless disseminate to all is selected, only clients in this list will receive the route being configured.

Alias/Group List. This list provides a list of all aliases and groups to which this route is currently not disseminated. Entries can be moved from this list to the dissemination list (to add it to the dissemination list) or vice versa (to remove it from the dissemination list).

(1) Route Management Embodiment

FIG. 12 illustrates the pre-call processing events of the example voice of internet communications system of the present invention. While all of the steps included here are not required or essential, they are illustrative of the features of the present invention. Routing Manager 352 maintains a list of local routes (step 1210). Other Routing Managers also maintain their own lists of local routes (step 1212). The Routing Manager connects to the routing servers for which is is configured (step 1218) and exports local routes which are enabled for export (step 1222). This same process is being performed by other Routing Managers (steps 1220 and 1228).

The exported routes are transformed on the routing servers (steps 1224, 1226, 1232, 1234) as described in FIG. 8. These exported local routes are then stored on the routing servers as dynamic routes. As discussed above in FIG. 9, disseminated routes are imported to the gateway servers from the routing servers (steps 1236 and 1242). For practical purposes, the gateway servers are now ready to handle calls and simply wait for them to come in (steps 1246 and 1248).

In response to call related information, the Routing Manager would make a query for a route (step 1250). Depending on the configuration of the system, various routing models can be used to provide potential routes. As discussed previously, dynamic and static global routes are cached (step 1254). The matching routes are prioritized (step 1258). Similar operations are being performed on other Routing Managers: In response to call related information, other Routing Managers would make a query for a route (step 1252). Depending on the configuration of the system, various routing models can be used to provide

potential routes. As discussed previously, dynamic and static global routes are cached (step 1256). The matching routes are prioritized (step 1260).

3. *Database Basics*

This section introduces the basic principles of operation behind the database server, as well as how gateway servers access the database server. This section provides the background knowledge necessary to understand and configure the services that depend on the database server, which are discussed in the following sections.

This section is divided into two sections, as follows. The first section deals with database issues on the gateway server side. The second with general issues on the database server side.

Gateway Access to Database Servers. Describes how a gateway server can be configured to access multiple database servers, and how multiple database servers can be used with each of the services provided via the database servers.

Database Server General Setup. Explains the use of the management system to modify general database server settings.

Database Server Access Control. Describes how to configure the clients of the database server, a necessary step in order to allow gateway servers and provisioning systems to connect to and make use of the database server.

1. *Gateway Access to Database Servers*

Similar to the Routing Manager component of the gateway server, discussed in the previous section, a Database Manager component also exists as part of the gateway server. This component is responsible for managing access to database servers, and for coordinating requests issued to multiple database servers. Double clicking on the "Database Server Access" root level node of the configuration tree accesses the Database Manager configuration screen 12602.

These directions should be performed when connected to a gateway server, not when connected to a database server. FIG. 126 shows the Database Manager configuration screen 12602.

5 The database server list is like the routing server list. It is a list of the database servers that the gateway server has been configured to access. When adding or modifying an entry in the database server list, dialog 12702, represented by FIG. 127, will be displayed, allowing one to enter information about the server.

10 *Server Name.* Name of the server being configured. This name will be used to identify this server entry when configuring the services that make use of database servers. Once the server has been added, the name assigned to it cannot be changed.

15 *Host Name/Port Number.* The hostname or IP address, and port number, of the database server that is being configured. The port number should be left at the default value unless you are specifically operating a database server on a different port number.

20 The Database Manager component allows Database services, described in the following two sections, to access any of the servers that are available. The configuration of how servers are accessed is similar to that for routing. Like routing, ordering information can be supplied that determines what order the servers are used in. database servers are organized into groups, much like routing servers are placed into ordering levels. Operations are performed on every database server in group # 1. If none of them handle the operation, then the operation is retried on every database server in group # 2. This continues until
25 no more groups are defined, or until at least one database server processes the operation successfully. There are two main differences, however, between database server request processing and routing server request processing:

A database server can be in more than one group. For example, Group # 1 could contain Server A and Server C, and Group # 2 could contain Server B and Server C.

Ordering/Grouping information is configured on a service-by-service basis. For example, the User Management service could use a different set of servers than the CDR Collection service.

2. Database Server General Setup

The ExpresConfig program was used to perform the initial set up of the database server. In the event that a change to the configuration of the database server is necessary, the management system should be used to make the change. Using the management system to administer such items as the management access control list, or identification/authentication configuration was discussed in above. However, it may also be necessary to make changes to the data source configuration that controls how the database server establishes a connection to the underlying Database.

When connected to a database server using the management system, double clicking on the "Database" root level node will display the main Database Server Window 12802, as shown in FIG. 128. By clicking on a "Change Settings" button within the Server Settings Panel 12804, it is possible to modify the general server settings that define the data source and database settings. Note that one should not change the port number under normal circumstances. The list of Clients (gateway servers) is shown in panel 12806.

3. Database Server Access Control

The database server manages a significant amount of critical data. User records and call records are essential to the correct operation of the system, and significant possibilities for fraud exist if an unauthorized user is able to obtain

access to modify or even just to view the data that is managed by the database server. For this reason, it is necessary to control access to the database server.

Database server access control works in the same way as management access control, and routing server access control. Like the routing server, it is necessary to create a client entry for each gateway server that connects to the database server. It is also necessary to create a client entry for every provisioning system that accesses the database server. The client list can be modified through the bottom portion of the main database window.

FIG. 129 shows the Client Configuration dialog 12902 of a single database client. The configuration of a single database client is significantly simpler than for routing clients since groups and aliases are not required in setting up the database server. General panel 12904 allows for the naming of a client as well as a short description. Panel 12906 displays the identities currently in use by this client.

4. *User Management*

This section covers the management of end users of the present invention, and describes the aspects of the system that are specific to handling details associated with the end users of the present invention. All aspects of user account setup, grouping of user accounts, accounting for calls made by users, setting up of debit information for users, and other such information is discussed in this section.

This section is divided into several sections that deal with different aspects of user management. The sections are as follows.

User Management Concepts. Discusses the main ideas of user management in the present invention. The main types of records that are maintained for the purpose of user management are discussed, and the interaction between these records is explained.

Gateway Server Access to Users. Describes the user management service that is part of the gateway server, and how it needs to be set up to obtain user records from the database server.

5 *Managing User Records.* Describes how the provisioning system can be used to set up and manage user accounts.

Managing Group Records. Shows how to use the provisioning system to set up and manage group records.

Managing Account Records. Demonstrates the setup and management of account records, including debit capability.

10 1. *User Management Concepts*

15 This section begins with a description of the fundamental purpose behind user management. Most importantly, user management is not necessary for certain applications. Most often, the purpose of user management is to allow an operate to charge for services provided by the present invention. For many enterprise applications, the identity of the user that is making a call using the present invention is unimportant, as is collecting records of all calls made by particular users. However, in providing a revenue generating commercial service, user management is critical.

20 User management provides several key things:

Identification and authentication of users. Without management of users, it would not be possible to distinguish between legitimate users, and unauthorized users of the system. Sometimes, the telephone equipment connected to the gateway server accomplishes this, but generally, user management is required in order to perform user verification.

25 *Classes of service.* In order to offer different classes of service (differentiated by numbers that can be called, or the quality of routes used to call those numbers), it is necessary to differentiate between users. User management

is required in order to do this. For example, it may not be desirable to provide all users with the capability to call internationally.

User-specific behavior. Some preferences and configuration may vary from one user to another. An example is the language preference of the user. This parameter may differ from user to user, and hence, it is necessary to differentiate between users.

Billing and Accounting. In applications where it is necessary to charge for services, it is necessary to identify the users that should be charged for the service, unless there is only one user of the system (which is usually not the case).

Debit Capability. When the debit capability of the gateway server is being used, it is necessary to manage users and their balances.

In the present invention system, all records associated with user management are stored on the database server. Other components of the system, such as the gateway server, must communicate with the database server in order to obtain those user records. This simplifies control over access to the underlying database, and allows multiple gateway servers to share the user database. Of course, since multiple database servers can be operated, it is not necessary for all gateway servers to use only a single database of users.

a. User Records

In order to provide the services described above, it is necessary to maintain a record of each legitimate user of the system. With that record, it is necessary to store information, such as their language preference. A user record is the piece of data that contains most of the information that is relevant to the user. It contains things such as:

- General information about the user, such as their name;
- The ID and password that a user requires to access the system;
- Additional identifying information, such as Caller ID, that can be used to automatically identify or authenticate the user;

- Preferences of the user that customize system behavior,
- The access level of the user to system services;
- Limited billing and accounting settings for the user.

An important piece of information in the above list is that only limited billing and accounting settings are stored in the user record. Most billing and accounting related settings are actually stored in account records, which may be associated with multiple users. The purpose of separating accounting information into a separate account record, described below, is to allow a single account to be shared by a group of users. For example, a corporate calling account might allow a group of users to be billed jointly, while still providing individual users with their own IDs and passwords.

b. Group Records

Group records are closely related to user records. The purpose of group records is to allow classes of users to be defined. A user may be part of either no group, or one group. If a user is part of a group, then certain settings in the user record may be left undefined, and specified instead by the group record. By taking this approach, rather than specifying all information directly in the user record, it is possible to make changes that affect a group of users simultaneously. For example, consider issuing a batch of 2000 demonstration calling cards. If these calling cards need to have an expiration date, and each calling card had a separate user account associated with it, then it would be necessary to specify the expiration date for each one of those calling cards. If it was desired to change the expiration date for that entire batch of calling cards, then every record would need to be updated, or specialized tools for provisioning the database would need to be used. Instead, it is possible to create a group record, associated with the user record for each card, and to make a single change to the group record.

For certain fields in the user record, it is possible to specify that the value in the associated group record should be used as the value of that field. For

example, the expiration date in a user record could be set to "use group record settings". With this setting, the group record would determine what the expiration date of the account was. These settings can be made on a field-by-field, user-by-user basis. For example, consider two user records associated with the same group record. User record # 1 could specify an expiration date of 01/01/2002, but with privilege levels defined by the group record. User record # 2 could have a group specified expiration date, but with a privilege level of 2. This provides the flexibility of using group records to represent default settings, while permitting specific changes and exceptions on a user-by-user basis.

Much of the information in a user record can be left unspecified, and determined by the values in the corresponding group record. These fields include, but are not limited to:

- Start and expiration dates for the account;
- Privilege levels, both normal and authenticated;
- Preferences, such as authentication requirement and language settings;
- Time of day usage restrictions;
- Billing and accounting information.

In general, because the use of group records in no way restricts the customization of each individual user record, it is recommended that group records are used to specify options that are common across a group of users.

c. Account Records

Account records provided the basis for real-time debit in the present invention system. Although account records can be used in post-paid applications, it is often not necessary to do so. This is because a third party billing application that prepares bills using the call detail records is provided with data that identifies the exact user that made a particular call. From this point, the third party billing application could use its own rules regarding whether the call

should be billed to the user, to a shared account, directly to a credit card, or in some other manner.

However, for real-time billing using the debit capability of the software, it is necessary to define an account record. Account records define properties that are used for billing, and may be associated with credit records, to limit the amount of credit extended to a user, or with debit records, which are used to provide real-time billing. Credit and debit information maintained by the system is always associated with an account record.

d. Remote Gateways and User Management

It is desirable to apply many of the capabilities provided by user management to gateway servers as well as to end users. For example, it might be desirable to have a debit account for a particular gateway server, or to control the permission levels granted to different gateway servers. In order to facilitate this, the gateway server is capable of associating a user account with remote gateway servers that request calls.

The mechanism for performing this association is by using the Remote System ID, a configurable parameter on each gateway server. This ID is used as if it was Caller ID information, and is used to look up a matching user record. Details on setting this up are described below, in the section on managing user records. It is also possible to use authentication information in the user record to verify the IP address of the remote gateway server.

2. Gateway Server Access to Users

This section describes how to set up the user management service on the gateway server. As described in the previous section, database services each supply an independent list of the database servers that are to be used, grouped according to the order in which they should be used, to obtain user records. For gateway servers performing user management, the principle type of operation

used is that of obtaining a user record from the database server. It is generally not necessary for this kind of operation to perform the operation on multiple database servers at the same time. However, it may be desirable to fallback to an alternate database server to obtain the user record if the normal database server that provides user records are obtained is not available. This section describes how to perform such a configuration of the user management service on the gateway server.

The server access settings for the user management service can be viewed by double clicking the "User Management Service" node, visible when the "Database Server Access" root level node is expanded. The Server Settings window 13002 will be displayed, as shown in FIG. 130.

In this window 13002, it is possible to see the definition of groups, and the members (i.e. database servers) that are members of each group. For the reasons explained above, there should normally be one database server configured in each group. The primary server from which user records are obtained should be in Group #1. The secondary server should be in Group # 2, and so forth. It is possible to add the name of a database server that is in a particular group by selecting the Group, then clicking on Add Member. The names entered for database servers must be selected from the list of configured database servers.

3. *Managing User Records*

The provisioning system is used to manage all user records. In order to use the provisioning system to access the user records managed by a database server, it is necessary to ensure that one has a user management session open in the provisioning system. Once a user management session is open, one is presented with Query screen 13102, shown in FIG. 131.

From screen 13102, one is able to use Query Options panel 13104 to query the database of existing users, add new users, remove users, and modify existing users. Note that there are two buttons: "User Properties," which is used

to get details about a particular user, and "Modify," which is used to change an existing user record.

When adding or modifying a user record, one will be provided with a user information dialog that consists of a number of panels. This is discussed in the subsections below. However, prior to modifying a user record, you will need to query for the records that you wish to modify. This is also required if one wishes to remove a user record. First, using the interface shown above to query existing user records is discussed.

a. User Record Queries

The Query screen 13102 is the main interface to user management capabilities on the provisioning system. Screen 13102 allows user records to be searched, based on various criteria.

In order to query available user records, one or more query options can be specified. User records can be queried by a variety of fields, including name, ID, and authentication code. When specifying the search criteria, it is possible to use "*" as a wildcard that will match any string. For example, entering "12*" as the User ID search criteria will find all user records beginning with 12. Once the search criteria are entered, the ordering parameter for results should be selected. This controls the order in which the results are displayed. Once this parameter has been set, clicking on Query will begin the search for matching user records.

Once the query is complete, the results will be displayed, and the "Close Query" button will become selectable. The "Close Query" will close the query, and allow a new query to begin. Once a query is closed, it is no longer possible to travel forward and backward through the results. However, any user records displayed in the result window can still be selected, and viewed or edited. If there are more user records matching the query than can be displayed at a time, the "Next" and "Previous" buttons will allow scrolling through the matching user records.

Group/Account Record panel 13106 allows for the review of group and account records related to a queried user. User list window 13108 displays a list of found users based on a query entered and performed in panel 13104.

Whether or not a query is active, it is possible to add a user by clicking on the "Add User" button. However, the other operations require that a user record in the result window of a query be selected. Once selected, the user record can be removed, viewed, or modified.

b. User Record Information

FIGs. 132 - 137 show the configuration options for a new user. There are six panels that are used to display information about a user. These panels are shown in Tabs 13204 of FIG. 132. Each one of these panels is discussed below.

(1) Identification Panel

Identification panel 13206 contains information that is used to identify and authenticate the user, as well as some general settings for the user account, as shown in FIG. 132 and described below.

Name. Name of the user. This field is used for information purposes. The name field is optional, but once a user has been added, the name for that record cannot be modified.

Numeric ID. This is the ID of the user. Since users generally access the system using a telephone, it is required that the ID only contain numbers that can be entered using DTMF tones. The ID of the user must be unique.

Authentication. This group of fields specify whether it is automatically necessary for the user to specify their authentication code. Before a user enters their authentication code, they will be assigned their normal privilege level. Once their authentication code has been entered, they will be granted their authenticated privilege level. If a user is required to always enter their authentication code, they will be prompted to do so regardless of whether or not

the service that they are accessing actually requires that privilege level. Otherwise, if they are not required to always enter their authentication code, they will only be asked for their authentication code if it is necessary in order to make use of a service that they have requested access to. It is possible to use the settings defined in the associated group record for this field by choosing "Use Group".

Language Preference. Specifies the language preference for the user; if left blank, the default language will be used.

E-mail Address. Used to record the E-mail address of the user. Available for information purposes only, and can be left blank.

Home Phone. Home telephone number of the user. Available for information purposes only, and can be left blank.

Online Announcements. Specifies which announcements the user wishes to hear when making debit calls. If "debit balance" is selected, then the user will be told their remaining balance each time they place a call. If "time left" is selected, then the user will be told the maximum duration of a call each time a call is placed. If "use group" is selected, then debit balance/time left prompting will be performed based on the group record setting for the value.

(2) Permissions Panel

Permissions panel 13302 is used to configure the access level, and account active dates for the user account. Permission panel 13302 is shown in FIG. 133 and described below:

Account Status. This field allows an account to be enabled or disabled. Group settings can also be used to determine whether or not an account is enabled.

Permissions. Determines the normal and authenticated privileges assigned to this user account.

Activation Date/Time. Sets the date and time when the account will become active. The account may only be used on or after the time specified. It is possible to use the group value for this setting.

Expiration Date/Time. Sets the date and time when the account will expire, and become unusable. The account may not be used after the specified time. It is possible to use the group value for this setting.

(3) *Caller ID Identification Panel*

Caller ID Identification panel 13402, shown in FIG. 134, allows additional identifying information about the user to be entered. When a user accesses the system, identifying information may be available. Examples include Caller ID information signaled by the telephone equipment, the H.323 terminal name of an H.323 user, or the remote system name of a calling gateway server. Any of this information can be used to automatically identify a user, without requiring them to enter their ID. This panel, shown in FIG. 134, allows you to configure the information that can be used to automatically identify a user.

In order to add new identifying information for a user, type the value associated with that information in the New String field. For example, if the Caller ID string that should be used to identify the user is 7035551234, then enter this string in the New String field. Once the string has been entered, click Add. A new entry in the table, with the corresponding string, will be added. Once it has been added, click on the lookup type column for that string. One will be able to select what type of identification the new string corresponds to. The available options are:

- Caller ID / ANI. The entered string will be matched against Caller ID information provided by the connected telephone equipment.
- H.323 terminal name. The entered string will be matched against the H.323 terminal name supplied by the user when connecting using the H.323 protocol.

- Remote system name. The entered string will be matched against the system name supplied by a remote gateway server.

In order to remove an identifying string, simply select it and click Remove.

Note that identifying information must be unique; in other words, the same Caller ID string and type cannot be associated with more than one user record. This is clearly necessary since the information is used to uniquely identify a particular user.

(4) Caller ID Authentication Panel

Caller ID Authentication panel 13502, shown in FIG. 135, is the analog of identification panel 13402, but applies to authentication of the user. It is configured in a manner that is very similar to the identification panel discussed above.

As with the identification panel, it is necessary to type in the authentication string, and click on Add in order to add an authentication record with that string. Then, click on the lookup type column to specify what the lookup type for that string should be. There are numerous authentication types, some of which are for expansion purposes.

- Caller ID/ANI. This allows caller ID information to be used to authenticate the user. This allows a single Caller ID string to be associated with multiple users, and allows Caller ID to bypass password checking. For example, it is possible to set up a calling card account that when used from the users home, requires only the users ID to be entered, but when used from a payphone, requires both the user ID and authentication code.

- IP address and port. Must be entered in "192.168.1.3 port 6789" format, and will authenticate the user only if their IP address and port number match the specified values.

- IP address only. Must be entered in “192.168.1.3” format, and will authenticate the user only if their IP address matches the specified IP address. IP address with netmask. Not currently used for authentication purposes.

- PIN code. This authentication type is not presently used, since the authentication code is contained in the basic user record settings.

- Text password. Not currently used for authentication purposes, designed for enhanced applications in the future.

Note that the reason that there are some additional types is that third party applications could make use of these types even though they are not directly used a gateway application.

As with the identification panel, authentication entries can be removed by selecting the entry and clicking on “Remove”.

(5) Account Record Status Panel

The account record status panel 13602 is used to display information about the account record that is associated with this user record. It allows an account record to be selected or created, and is shown in FIG. 136 and described below.

Clicking “Select New” will allow the selection of an account record that is associated with the current user record. It is possible to select either an existing account record, or to create a new account record associated with the user record. When select new is selected, the account record query window will be displayed, allowing existing account records to be searched. This process is described in greater detail when account records are discussed, below.

The “Change Settings” button allows the account record associated with the user record to be modified.

The “Detach Account” button will remove the association between the account record and the user record.

(6) Parent Group Panel

The parent group panel 13702, shown in FIG. 137, displays information about the group associated with this user record, if any, and allows the group association to be changed.

5 Clicking "Select New" will allow the selection of a group record that is associated with the current user record. It is possible to either select an existing group record, or to create a new group record. This is done by way of the group record query window, which will allow you to search for the group record that you wish to add. The group record query screen is discussed in greater detail below.

10 The "Detach Group" will remove the association between the user account and the group record with which the user account is currently associated.

4. Managing Group Records

15 The management of group records is very similar to the management of user records, since group records mainly provide default values for user record fields. The group record query screen 13802, from which group record operations can be performed, is accessed via the user record query screen 13102, by clicking on the "Group Records" button in panel 13106. The group record query screen 13802 is discussed in the sections that follow and shown in FIG. 20 138.

a. Group Record Queries

25 Through Query Options panel 13804, queries of the available group records on the system is done. Querying works very similarly to querying user records. The first step is to specify the search criteria, by entering information into the fields on the left that correspond to the group records that you wish to retrieve. As with user records, it is possible to use wildcards when searching for

group records. Ordering parameters can also be set, following which it is possible to query for group records.

Once results have been returned in a list within Account Records panel 13806, group records work identically to user records; it is possible to close the query, scroll forward and backwards using the “Next” and “Previous” buttons respectively, and to add/remove/view/modify group records.

b. Group Record Information

The information available in the group record is a simplified set of the information contained in the user record, since not all user record fields can be left as defined by the group record. There are two panels that provide access to the fields contained in the group record, each of which is described below.

(1) General Panel

Most settings are available on general panel 13906, shown in screen 13902 of FIG. 139. For an explanation of any of the parameters shown, refer to the corresponding parameter in the user record, where the parameter is actually described. User records will only use values specified in the group records if the corresponding user record values indicate that the group record should be used to determine the value.

The account ID setting works identically to that for the user record, but will establish an association between the group record and the account record.

The only group record specific information is the “Group Info” field, which allows descriptive information about the group to be entered. This information is for descriptive purposes only, and does not affect the behavior of the system.

5. Managing Account Records

Management of account records can be performed through the account record query screen 14002, accessible by clicking on “Account Records” when viewing the user record query screen 13106 in the provisioning system. This brings up the account record query window 14002, shown in FIG. 140.

a. Account Record Queries

The account record query window 14002 functions in the same manner as for user records and group records. The first step is to input query information in the fields in query panel 14004, also select ordering settings, and click on the “Query” button. Like user and group records, it is possible to scroll through found account records shown in Account Records panel 14006 using the “Previous” and “Next” buttons.

The “Add New” button allows new records to be added; once a query result has been selected, the “Remove”, “Properties”, and “Modify” allow the account record to be deleted, viewed, or modified, respectively.

b. Account Record Information

Account record information is administered in panels 14106 and 14202. These are shown in FIGs. 141 and 142 and described below.

(1) General Panel

General panel 14106, shown in FIG. 141, allows general account record information to be configured. Tabs 14104 allow for switching between General Panel 14106 and Profiles and Restrictions panel 14202. An explanation of the fields of panel 14106 follows.

Account ID. Specifies the ID of the account. This ID is referenced by user records and/or group records, and must be unique. If account records are

5

10

15

20

The profiles and restrictions panel 14202, shown in FIG. 142, is used to control the rating of calls, and modes of operation which are supported on this account.

25

Enable Credit. Enables credit to be used on this account.

Enable Billing Exempted. If set, this flag exempts the account record from being charged for any use of services.

5. *CDR Collection*

This section presents call detail record (CDR) collection, and explains how this task is accomplished in the present invention.

CDR collection is essential in any environment in which charges are applied for the services provided by the present invention . A gateway server will create a CDR for every call that is placed through the system. CDR collection is the process of collecting the CDRs across multiple gateway servers, and storing them in a centralized database. As such, CDR collection involves the database server as well as the gateway server.

The discussion of CDR collection is divided into three main sections:

Querying CDRs on a Database Server. Explains how the provisioning system is used to access the CDRs that are available on a database server.

Information in a CDR. Describes the information in a CDR, and the meaning of that information.

Gateway Server Collection of CDRs. Describes the process by which the gateway server collects CDRs, and how it transmits these CDRs to a database server, or to multiple database servers.

Database Server Failure Handling. Describes how the gateway server handles inserting CDRs in light of a temporary failure of the database server to which the CDR is inserted.

1. Querying CDRs on a Database Server

Similar to the operation of user management, it is possible to use the provisioning system to access the CDRs that are contained in the database managed by a database server. This is not the only way to access CDR information; it is also possible to directly access the underlying database.

In order to query CDR information, it is necessary to establish a CDR session with the database server using the provisioning system. Once such a session has been established, the CDR query screen 14302 will be displayed, as shown in FIG. 143.

From window 14302, CDRs can be queried in a manner that is similar to that used for querying user, group, and account records, described in the previous section. Search criteria can be entered in the available fields of Query Options panel 14304. After search criteria is entered, the ordering parameter should be selected. Then clicking "Query" will display a list of CDRs that match the query in panel 14306. It is possible to use the "Next" and "Previous" buttons to scroll through the available CDRs. The "Close Query" button should be used if it is desired to make an other query.

The information contained in each CDR is described below.

2. Information in a CDR

CDRs collected by the gateway server contain a large amount of information. This section will help one to understand the fields in an example CDR. The present invention is not limited to this example CDR and information can be added or removed from the CDR depending upon a particular billing service or system requirements.

System Name. Contains the identity of the gateway server that submitted the CDR.

Inbound Intf. Identifies the inbound line group that received the call.

Inbound Line. Identifies the specific line on a line group that received the call.

Outbound Intf. Identifies the outbound line group used to place the outbound portion of the call. May be unknown if the call was never placed.

5 *Outbound Line.* Identifies the specific line on the outbound line group that was used to place the call.

Call Type. Indicates the charge type for the call.

Called Addr. Indicates the address that was called. This field will generally be in E.164 form, but is prefixed with the "e164:" tag to identify E.164 numbers.

10 *User Connect Time.* Indicates the time, as collected by the submitting gateway server, at which the user connected to the system.

Call Connect Time. Indicates the time, as collected by the submitting gateway server, at which the call placed by the user was connected to the destination party.

15 *Call Duration.* Specifies the duration, in seconds, of the call. This is the time between the called party answering the call, until the time that the call was disconnected.

IVR Duration. Specifies the amount of time, in milliseconds, that the user spent in the IVR process, entering their ID, the number to call, and so forth.

20 *Setup Duration.* Specifies the amount of time, in milliseconds, required to set up the call. This includes searching for routes, obtaining a line to place the call, contacting a remote Gateway, and other overhead.

Dial Duration. Specifies the amount of time, in milliseconds, that the call was placed for before it was either terminated, or until the called party answered.

25 *Call Attempt.* Depending on the configuration of the gateway server, it may be possible to make multiple calls on a single call into the system. This field indicates the call attempt number within a single call into the system.

Identified. Indicates whether or not the user making the call was identified.

Succeed. Provides a success code indicating whether or not the call was successful, and if the call was unsuccessful, provides a cause code.

5 *RemoteSystem.* Indicates whether the call was originated by a remote gateway server, or a local user of the system.

User Name. Provides the name, if available, of the user that placed the call.

User ID. Provides the ID, if available, of the user that placed the call.

10 *Caller ID.* Stores caller ID information, if available, of the caller that accessed the system.

Group ID. Stores the group ID, if available, of the group account associated with the user when the call was made.

15 *Account ID.* Specifies the account ID, if available, associated with the user at the time that the call was made.

Session ID. Provides a unique call ID that identifies this call. The session ID can be used to correlate the CDR collected by the originating gateway server, and the terminating gateway server.

20 *Route Used.* Provides information, if available, regarding the route used to make the call.

Inbound Info. Provides information available regarding the inbound line or user.

Outbound Info. Provides information available regarding the outbound line or user.

3. Gateway Server Collection of CDRs

CDR collection is a capability that is provided by the gateway server platform. The gateway application makes use of this capability to create a CDR for every call that it processes. CDRs are inserted regardless of whether or not the call was successful; the success code, explained above, is used to record the actual success of the call.

On the gateway server side, the configuration of CDR collection involves specifying the database servers list and groups that are used in CDR collection. Double clicking on the "CDR Service" node can access the CDR collection configuration dialog. This node is visible when the "Database Access" root level node in the configuration tree is expanded, and is shown in CDR Service window 14402 of FIG. 144. Note that one should use the management system to configure these parameters, and that one should be connected to a gateway server in order to do so.

The server settings for the collection of CDRs is configured as described above. For the CDR service specifically, it is useful to understand how concurrency and fallback work. If two or more database servers are placed in the same group, then both of them will receive a copy of the CDR when it is inserted. Only if all database servers in that group fail to receive a copy of the CDR will the next group be used. This can be used to achieve redundancy, by always logging CDRs to multiple database servers. (However, there is still some resilience to database server failures even in the absence of such a configuration, as described in the next section). Configuring independent groups, allows a secondary database server to receive CDRs in the event that the normal database server is not available.

4. *Database Server Failure Handling*

In the event that none of the database servers are available to log a CDR collected by a gateway server, it remains important that the CDR not be lost, as this would represent lost revenue. Therefore, for CDRs that cannot be successfully transmitted to a database server, a local queue is maintained. While the database server remains unavailable, any CDRs collected will be stored in this queue. Once the database server becomes available, the CDRs in the queue will be transmitted to the database server, and then removed from the queue. This queue is maintained persistently on disk, so that even if a gateway server undergoes a temporary failure while the database server is unavailable, it will continue to try to insert CDRs once it recovers from the temporary failure.

6. *License Server*

The Licensing Server is an additional (optional) software component according to the present invention. The License Server is designed to simplify the administration of a VoIP system according to the present invention, by allowing all licenses to be managed from a single centralized server. This eliminates the need to configure license keys manually on individual components. Furthermore, it allows licenses to be purchased in large blocks and then distributed amongst gateway servers, rather than needing to purchase large numbers of small quantities of licenses for each gateway server.

It is possible to operate a secondary License Server for backup a purpose, which is capable of allowing a network to continue to operate in the event that the primary license server becomes unavailable.

III. APPENDIX A

1. Database Table Descriptions

This appendix is intended to provide a description of the database tables used by the database server. The purpose of this description is to allow third party applications to be written which interface with the data maintained by the database server. Examples of such applications include billing systems, or customer account creation tools. These database tables are an examples of the present invention and do not necessarily limit the present invention. Different fields can be added or removed from the database tables as would be apparent to a person skilled in the relevant art given this description.

Each of the tables created in the database are described in the sections that follow. For information about how the database server uses the information in these tables, refer to the sections on User Management and CDR Collection.

1. ACCOUNT_RECORD Table

The account record table is used to store the list of accounts that are available for use on the system. For a description of the interaction between account records and user/group records, see the section on User Management.

a. Table Definition

Column Name	Description
Account_ID	Numeric ID representing this particular account type. (The primary key for the table should be unique.)
Account_Info	Account name, label or description.
Account_Rules	Account parameters
Account_Restrictn	Account limitations
Concurrent_Max	Maximum number of users/groups that can access this account type concurrently (i.e. at the same time)

Column Name	Description
Acct_Rating_Prof	Profile of Account's rating information.
User_ID	<p>This field will contain either one <u>specific</u> user ID or a "0" (zero). If there is only <u>one</u> user associated with this account ID, the corresponding user ID appears in this field. If there are <u>multiple or NO</u> users associated with this account ID, a "0" (zero) displays in this field. (The individual user IDs of users associated with this account ID because of their group association, will not display) <i>Examples:- If user ID 999 is the ONLY user associated with this account (user ID 999 is not associated with any group), "999" (without quotations) appears in this field. [Although there is only one individual user associated with the account, there may also be groups which are associated with the account as well – see group_ID below.]- If user ID 123 is a member of group ID 456, which is associated with this account (see GROUP_RECORD table), one would NOT enter user ID 123 (or any other member of group ID 456) in this field.- All users/groups associated with an Account ID, SHARE billing – if, in the above example, one wanted to provide INDIVIDUAL billing for user 123 (and all of the members of group ID 456), one would NOT associate the group (456) with any account, but would create a separate account ID for EACH user ID – in which case, in each account (ID) entry, this field would contain the associated user ID. - If there is more than one user ID associated with this account (who are not part of the same group), a 0 would display in this field to indicate multiple user IDs. - If there are NO individual user IDs associated with this account (again, this does not include users who are part of a group associated with this account), a 0 would display in this field to indicate NO individual user IDs.</i></p>

Column Name	Description
Group_ID	<p>This field will contain either one <u>specific</u> group ID or a “0” (zero). If there is only <u>one group</u> associated with this account ID, the corresponding group ID will appear in this field. If there are <u>multiple or NO groups</u> associated with this account ID, a “0” (zero) displays in this field. If a group is associated with an account, the corresponding account ID appears in the group’s entry in the GROUP_RECORD table. Once a group is associated with an account ID, all members (users) of the group automatically become linked to the same account ID.</p> <p><i>Examples:</i></p> <ul style="list-style-type: none"> - If only group ID 111 was associated with this account, “111” (without quotations) appears in this field. - If both group ID 111 and group ID 222 are associated with this account, a 0 displays in this field to indicate multiple group IDs. <p><i>Remember that ALL users/groups associated with an Account ID, SHARE billing. To provide separate billing for each group (all group members are “lumped together), create a separate account ID for EACH group ID – in which case, in each account (ID) entry, this field would contain the associated group ID. [To provide separate billing for each member of a group – instead of lumping them together – do NOT associate the group with an account. Instead, create a separate account ID for EACH user ID – see User_ID field above.]</i></p>
Time_Stamp	Leave blank – DB will enter information as required to ensure mutual exclusion on modifications.
Has_Debit_Record	A boolean type value where TRUE=account is using an (online) debit plan. FALSE=no associated debit balance.

Column Name	Description
Has_Credit_Record	A boolean type value where TRUE=account is using the credit billing plan. FALSE=not on credit plan. <i>(Speeds up internal DB operation – DB server looks at this field first to determine if it needs to fetch the credit record.)</i>
Locking_By	Leave blank – DB will enter information as required to ensure mutual exclusion on modifications.
Billing_Flags	Binary data indicating which billing method is used – debit (prepaid), credit (usage-based) or exempted (not charged). This field cannot be edited manually. If entering information DIRECTLY into the SQL/Access database, enter NULL.

b. Additional Notes

The “User_ID” and “Group_ID” fields are closely related. These two fields will contain either a specific user/group ID (respectively) or a “0” (zero). The value of zero (“0”) means that there isn’t one specific user/group ID associated with this account ID – there could be several users/groups associated with this account ID, or there could be none.

In one example, one of these fields (“User_ID” or “Group_ID”) should have a value of “0” (zero) at all times. This means one should not have one (individual) user AND one group associated with the account at the same time.

If a user is associated with a group ID which is, in turn, associated with a particular account ID, you do NOT have to enter the individual user ID in the “User_ID” field – the user will be associated with the account by group association. All users/groups associated with a particular account ID are under ONE single billing plan. To provide individual/separate user billing, do NOT associate the group with an account, instead create a separate account ID for each user ID. This does not mean you cannot still associate the user with the group – a user can be associated with a group but have its own account ID for billing

purposes. However, a user cannot be a member of a group which IS associated with a particular account ID and have its own account ID at the same time.

2. *AUTHENTICATE_ID Table*

5 The AUTHENTICATE_ID table is used to hold authentication information for both user records and group records. Since a particular user or group record may have multiple authentication entries associated with it, it is not possible to store this information within the user record or group record itself.

a. *Table Definition*

Column Name	Description
10 User_OR_Group_ID	The user ID <u>or</u> group ID which will be authenticated by the information contained in the "Authenticate_ID" field below. If you enter a group ID, all members of the group will <u>share</u> the <u>same</u> authentication information – each member still has their own numeric user ID. If one wishes to assign separate authentication information to each user (in a group), create a separate entry (in this table) for each user ID. A user can be a person, remote Gateway or H.323 user.

Authenticate_ID

Enter the information that will be used to authenticate the user or group. This information corresponds to the authentication method indicated in the "Lookup_Type" field below.

Examples:- If the user/group is authenticated by entering a numeric code such as 666, this field should contain "666" (no quotations).- If the user/group is authenticated by their ANI information, this field should contain the area code and phone number, no brackets or spaces, from which they are authorized to call, e.g. 4167360901.- If the user (Gateway) is authenticated by its IP address, this field should contain the IP address from which the remote Gateway will be allowed access. Enter the IP address as follows: 255.255.255.255- If the user (Gateway) is authenticated by its IP address plus port, this field should contain the IP address from which the remote Gateway will be allowed access and the port through which the communication will be delivered. Enter the IP address and port as follows: 255.255.255.255 port 0

Is_User_Entry

This field indicates whether or not this entry is for an user (as opposed to a group). A boolean value where "0" (zero) indicates this entry is for a user. "1" indicates this is a group entry.

Lookup_Type

This field contains an integer value identifying the authentication method to be used. For example, if this field contains a "0" (zero), the user or group specified in the above "User_OR_Group_ID" field will be authenticated by the ANI information contained in the above "Authenticate_ID" field.

b. Additional Notes

The "Lookup_Type" field values are as follows:

0 = User authentication based on numeric PIN/code.

1 = Authentication based on text password.

2 = User authentication based on ANI information.

3 = Gateway authentication based on IP address only.

4 = Gateway authentication based on IP address + port.

5 = Gateway authentication based on IP address + netmask.

3. CALL_RECORD Table

The call record table is used to store call records for a number of Gateways. These call records can be used for post-paid billing purposes, and allow custom billing applications to be developed which interface with the system. The CALL_RECORD table is updated continuously as calls are completed.

a. Table Definition

Column Name	Description
VOIPgate_System	Gateway writing the call record.
Session_ID_High	Unique identifier for the call that is consistent between CDRs inserted from both Gateways involved in the actual call. Can be used to correlate call records from the originating and terminating Gateways. Contains the high 32 bits of the 64 bit call ID.
Session_ID-Low	Same as above, but contains the low order 32 bits of the 64 bit call ID.
Inbound_Intf_Type	Interface used for inbound call (telephony or network).
Inbound_Inft_Inst	Instance number of the inbound interface.

Column Name	Description
Inbound_Line_Type	Hardware line type used for inbound call.
Inbound_Line_Inst	Instance number of the inbound hardware line.
Outbound_Inft_Type	Interface used for outbound call (telephony or network).
Outbound_Inft_Inst	Instance number of the outbound interface.
Outbound_Line_Type	Hardware line type used for outbound calls.
Outbound_Line_Inst	Instance number of the outbound hardware line
Call_Addr	Phone number dialed/called.
User_Connect_Time	Date & time PSTN user contacted the Gateway (to initiate a call).
Call_Connect_Time	Date & time call was connected (to called party).
Call_Duration (secs)	Duration (length) of the call (in seconds).
IVR_Duration (ms)	Length of time required to prompt user through the call.
Setup_Duration (ms)	Length of time required to make the call. [<i>This is the time required for the Gateway to answer the user's initial call (call connect), receive the number (entered by user), authenticate user, locate matching network call route, route the call via the IP network, have the second remote Gateway dial the number, until a "call connected" signal indicates call has connected to called party – "User_Connect_Time" to "Call_Connect_Time".</i>]
Dial_Duration (ms)	Time required (for second/remote Gateway) to dial outbound number and wait for call to connect.

Column Name	Description
Call_Attempt	Number of consecutive calls made by the user. <i>Example: When <u>called</u> party hangs up, the Gateway returns control to the (initiating) user, prompting the user for another number. If the user enters another number (to make a 2nd call), the call record for the first call will have Call_Attempt=1 and the 2nd call record will show Call_Attempt=2.</i>
Is_Identified	A boolean value reflecting whether a user ID was entered. TRUE= <u>no</u> user ID entered (typically this displays when call record is written by remote Gateway); 0=user ID was entered (typically this displays when call record is written by initiating Gateway).
Is_Succeed	Indicates whether call was successful or failed.
User_Name	User NAME. If the initiating Gateway writes the call record, the user name of the CALLER will display. If the remote Gateway writes the call record, the initiating Gateway's System Name will display. If the user/Gateway was not identified, "unknown" will display.
5 User_ID	User ID of the caller or Gateway (if applicable). If user/Gateway was not identified, "unknown" will display.
Caller_ID	ANI information of the caller (IF this service is available on the phone lines connected to the Gateway). This information only displays on call records from the initiating Gateway.
Route_Used	Identifies how the Gateway routed the call. If the Gateway routed call out over the <u>network</u> interface (to a remote Gateway), the IP address of the terminating Gateway will display here. If the Gateway dialed the number out over its phone lines (<u>telephony</u> interface), the phone number dialed will display here.

Debit_Reserved	For internal use only. Before a call is made, the database server reserves a number of credits for the call – reserved credit is not available for other calls. At the end of the call, the ACTUAL credit used is deducted from the account's debit balance and the remaining (unused) reserved credits are released – and are available for another call.
Debit_Sessions_Num	For internal use only. Indicates number of online debit sessions (under this debit ID) currently in progress.
Debit_Max_Negative	This defines the MAXIMUM credit that can be overdrawn from this debit ID balance. This field must have a value that is equal to or greater than zero. <i>For example, if this field has a value of 100, the system will (once the debit balance has been used up) allow \$100 worth of calls to be made (overdrawn) before disabling the account. If you do not want to allow ANY calls after the debit balance has been used up, enter "0" (zero).</i>
Debit_Reserved_Info	For internal use only. Contains binary information.

5

b. Additional Notes

10

15

This table – DEBIT_RECORD – and the ACCOUNT_RECORD table are linked by their "Debit_ID" and "Account_ID" fields respectively. An *account* is linked to the DEBIT_RECORD entry whose Debit_ID field has the SAME VALUE as that contained in its Account_ID field. To assign an online debit balance to an *account*, you create a new entry in the DEBIT_RECORD table and enter the *account* ID in the Debit_ID field. For example, to assign an online debit balance to *account* ID "1234", one creates a new entry in the DEBIT_RECORD table and enter "1234" in the Debit_ID field. This means that ALL members of account ID "1234" will use the online debit balance defined under the Debit ID "1234".

If entering online debit balance information directly into the SQL/Access database (instead of via the provisioning system), enter information in the Debit_ID, Debit_Balance and Debit_Max_Negative fields. The remaining fields should be left blank.

5 All members of an account can SHARE the SAME online debit balance
To assign a separate online debit balance to each user, create a separate account ID and corresponding debit entry for EACH user.

6. *EVENT_LOG Table*

The event log table is used to store events logged to the database server.

a. *Table Definition*

Column Name	Description
VOIPgate_System	Name of server writing the event log.
Event_Code	Type of event (general, resource warning, state change)
Event_Date	Date and time of the event.
Event_Info	Description or additional information.

7. *GROUP_RECORD Table*

The group record table allows users to be grouped, and to use default values for certain fields in the user records that are defined by the group record.

a. *Table Definition*

Column Name	Description
Group_Info	Group name/label.

Column Name	Description
Acct_First_Date	Date and time the group is activated (granted access).
Acct_Last_Date	Date and time group access is disabled/removed. Expiration date.
Concurrent_Max	Maximum number of users that can access the Gateway under this group name concurrently (i.e. at the same time). If this group is associated with an account, the "Concurrent_Max" field value in the ACCOUNT_RECORD table will be used.
Normal_Privilege	Privilege assigned to the user after they enter their numeric user ID (or, when they have been identified by an alternative identification method such as via ANI information).
Auth_Privilege	Privilege assigned to the user once they have been identified <u>AND</u> authenticated (by whatever methods the Gateway administrator has defined for the user). Normal and Authenticated privilege level can be the same.
Account_ID	If this <i>group</i> is associated with a particular <i>account</i> , the account ID will display here – if there is NO associated <i>account</i> , this field will contain a "0" (zero).
Time_Stamp	Leave blank – DB will enter information as required to ensure mutual exclusion on modifications.
Locking_By	Leave blank. The DB server sets and unsets this field according to usage – if account is being accessed via the provisioning system, the account will be locked.
Group_ID	The numeric ID representing this group.

Column Name	Description
Acct_Disabled	Through the provisioning system user interface, administrators can disable access for the group by checking the “disable” box in the Account Status section on the Permissions panel. This field contains a boolean value where “1” or TRUE=account disabled; “0” or FALSE=account <u>active</u> . <i>[If the group record is disabled, the Gateway will not allow calls from members of this group ID.]</i>
Force_Auth	Authentication is not always required – it is based on the required privilege of the desired route. Through the provisioning system user interface, administrators can force Gateways to always identify <u>and</u> authenticate users by simply checking a box. This field will display a boolean value where “1” or TRUE=box checked – user MUST be authenticated; “0” or FALSE=box not checked – authentication only when required based on privilege. <i>For example, if using calling card numbers as the user ID, one would want to force the user to also enter their PIN (i.e. authentication). “1” would display in this field.</i>
Restrict_Time	This field indicates whether or not, call time restrictions apply to this group, i.e. group members are restricted to making calls during certain time periods. This field contains a boolean value where “1” or TRUE=time restrictions apply – Gateway will check the “Allow_Times” field below; “0” or FALSE=no time restrictions – users can make calls at anytime.
Allow_Times	This field contains binary data specifying those times during which members of this group can make calls.

Column Name	Description
User_Flags	This field contains binary data. <i>[There are 64 generic flags defined (available for use) but only 2 flags have been assigned, in one example. The first bit specifies whether group members, when initiating a call, prefer to hear the <u>dollar amount</u> remaining in the online debit account. The second bit specifies whether group members, when initiating a call, prefer to hear how much <u>calling time</u> remains in their online debit account.]</i>

8. USER_RECORD Table

The USER_RECORD table is used to store user accounts. Users can actually be added to or removed from the system by modifying this table directly. However, it is important to take care to ensure that references to other tables (such as group and account records)

a. Table Definition

Column Name	Description
User_Name	Real (full) name of the user or company.
User_ID	Numeric user ID for this user – enter up to 15 digits. <i>[Users are identified and assigned privileges based on this numeric ID.]</i>
User_Auth	Numeric authentication code for the above user. <i>[Users may be required to enter this 2nd code to make a call.]</i>
User_Phone	Contact phone number for this account. <i>[This field is informational only – it is not used by the DB server.]</i>
User_Email	Contact email address for this account. <i>[This field is informational only – it is not used by the DB server.]</i>

Column Name	Description
Acct_Disabled	Through the provisioning system user interface, administrators can disable access for the user by checking the “disable” box in the Account Status section on the Permissions panel. This field contains a boolean value where “1” or TRUE=account disabled; “0” or FALSE=account <u>active</u> . <i>[If account is disabled, the Gateway will not allow calls using the specified user ID.]</i>
Force_Auth	Authentication is not always required – it is based on the required privilege of the desired route. Through the provisioning system user interface, administrators can force Gateways to always identify <u>and</u> authenticate users by simply checking a box. This field will display a boolean value where “1” or TRUE=box checked – user MUST be authenticated; “0” or FALSE=box not checked – authentication only when required based on privilege. <i>For example, if are using calling card numbers as the user ID, one would want to force the user to also enter their PIN (i.e. authentication). “1” would display in this field.</i>
Acct_First_Date	Date and time user ID is activated (granted access).
Acct_Last_Date	Date and time user access is disabled/removed.
5 Restrict_Time	This field indicates whether or not, call time restrictions apply to this user, i.e. user is restricted to making calls during certain time periods. This field contains a boolean value where “1” or TRUE=time restrictions apply – Gateway will check the “Allow_Times” field below; “0” or FALSE=no time restrictions – users can make calls at anytime.
Allow_Times	This field contains binary data specifying those times during which members of this group can make calls.

Column Name	Description
Concurrent_Max	Maximum number of users that can access the Gateway under this user ID concurrently (i.e. at the same time). If this user ID is associated with an account, the "Concurrent_Max" field value in the ACCOUNT_RECORD table will be used.
Normal_Privilege	Privilege assigned to the user after they enter their numeric user ID (or, when they have been identified by an alternative identification method such as via ANI information).
Auth_Privilege	Privilege assigned to the user once they have been identified <u>AND</u> authenticated (by whatever methods the Gateway administrator has defined for the user). Normal and Authenticated privilege level can be the same.
Time_Stamp	Leave blank – DB will enter information as required to ensure mutual exclusion on modifications
5 Locking_By	Leave blank. The DB server sets and unsets this field according to usage – if account is being accessed via the provisioning system, the account will be locked.
Group_ID	If this <i>user</i> is associated with a particular <i>group</i> , the group ID will display here – if there is NO associated <i>group</i> type, this field will contain a zero (0).
Account_ID	If this user is associated with a particular <i>account</i> , the account ID will display here – if there is NO associated <i>account</i> type, this field will contain a zero (0).

Column Name	Description
User_Flags	This field contains binary data. <i>[There are 64 generic flags defined (available for use) but only 2 flags have been assigned, in one example. The first bit specifies whether the user, when initiating a call, prefers to hear the <u>dollar amount</u> remaining in their online debit account. The second bit specifies whether the user, when initiating a call, prefers to hear how much <u>calling time</u> they have left in their online debit account.]</i>
User_Lang_Pref	This field indicates the language of choice for online debit prompt, e.g. "English".

IV. APPENDIX B

1. Network Design Information

5 This section provides information to assist in designing and organizing a network in order to improve performance. It also provides some information and suggestions with regard to how the components of the present invention s could be organized.

1. Network Requirements

10 Three main components are used in the present invention: the gateway server, database server, and routing server. This section explains the network requirements of each of the components, so that the appropriate network bandwidth and performance can be allocated to achieve the desired degree of performance. Each component, and the type of communication used by that
15 component, is described separately.

a. Gateway Server

The gateway server runs on UDP/IP port 6789 for control data. Additional UDP/IP ports are used for media data. These additional ports can be allocated anywhere within a specified range. By default, the range of ports that can be used is 1025 to 65535. However, the port range can be specified to match a user's requirements.

The main kind of communication made by the gateway server is communication to other gateway servers, for the purpose of completing a call. A small amount of control traffic is exchanged to set up the call. After this point, actual media data is exchanged throughout the call. Because the present invention employs silence compression to reduce the required bandwidth during periods of silence, the media data is not a continuous stream. However, with the default frame size of 60ms per packet, one 66 byte frame will be sent every 60ms during active conversation. During silence, no data will be transmitted. In addition, a UDP/IP header of 28 bytes will also be present on each transmitted frame. Overall, this results in a data rate of 94 bytes per 60ms, or 1,557 bytes per second. Expressed in bits per second, this is 12,533 bits per second, or 12.5kbps. In addition, the underlying network (frame relay, ATM, etc.) may add additional framing data, further increasing the actual data rate that must be allowed for on the network. Note that this is the transmission capability required from the sender to the receiver. When both the sender and the receiver are talking simultaneously, this much bandwidth must be available in both directions. Often, network bandwidth is symmetric. A 33.6kbps modem can carry 33.6kbps of data in both directions simultaneously; similarly, a T-1 can carry 1.544Mbit/s in both directions at the same time. However, in the case that asymmetric bandwidth is available, it is important to be aware that the bandwidth requirements of the system are bi-directional.

Another critically important aspect of network performance on the communication between two gateway servers is that of latency, and jitter.

Latency, the delay introduced on packets from one gateway server to another, will be perceived by the end users of the system as a delay in the response of a remote user. Some latency is introduced by the gateway server itself in processing and decoding the media data. However, network latency should be kept to a minimum to ensure that the conversation between end users is a natural one.

Even more serious than latency is the problem of jitter. Jitter is the variation in latency from one packet to another. For example, if packets always arrived 50ms after being transmitted, then there would be a 50ms latency and no jitter. As another example however, if packet # 1 arrived 100ms after transmission, packet # 2 arrived 50ms after transmission, and packet # 3 arrived 150ms after transmission, there would be an average jitter of +/- 33ms. The reason that jitter is more critical than latency is because it can cause a packet to arrive too late to be useful. The effect is that the packet may be delayed enough that the end user will hear a pause in the voice that is talking to them, which is very unnatural if it occurs during the middle of a word or sentence. Jitter typically occurs when the network utilization is too high, and packets are being queued, causing delivery times to become unpredictable. The Internet, because of its complex structure, is often subject to varying degrees of jitter. The present invention software compensates for network jitter. Furthermore, handling jitter comes at the expense of latency, since the only way to handle jitter is to buffer additional data so that when the data arrives exceptionally late, continuous playback to the end user can be maintained.

In addition to gateway server to gateway server traffic, the gateway server also communicates with the routing and database servers, and is managed by the management system. Network factors related to these components are discussed below.

b. Routing Server

The routing server is involved in routing calls between gateway servers. However, the routing server is optional; routing information can be configured statically on individual Gateways.

When routing server is used, there are two modes that it can be used in. Depending on which mode is used, the networking requirements between Routing and gateway servers will vary. Routing modes are explained in detail in the call routing section, however, there are two basic modes of operation from a network standpoint

Served-based Static Routing and Dynamic Routing. In these modes, each time a Gateway need to make a call, it will query the routing server for a route. This occurs once, at the beginning of a conversation. Routes are typically quite small, and do not require additional bandwidth to be allocated (since the bandwidth during the call will be higher than during the call routing stage). However, it is important nonetheless since network latency will add to the call setup time, and the end user will be waiting for the system to initiate the call. Normally, this latency is not a problem. However, if the latency is on the order of 500ms, and call setup time is critical, then this may be a consideration. An important aspect of dynamic routing is the caching of routes. With dynamic route caching, not all calls will require a query to the server. However, there will still be calls which require the server to be queried, and these calls will incur the latency penalty.

Disseminated Routing. With disseminated routing, the routing server distributes the necessary routing data to each gateway server when it firsts connects. This is typically when the system is starting up. In this model, the Gateway will use a large amount of bandwidth during the initial connection, as the routes are disseminated from the server. However, after this point, there will be no subsequent queries, and no latency will be added on a per-call basis.

The selection of a routing model will be explained in more detail in the section on call routing. However, it is useful to keep the network requirements in mind when considering what routing model to use.

c. Database Server

There are two principal kinds of network traffic that interact with the database server.

The first kind of traffic is due to communication with the gateway servers. There are two reasons for this traffic; user management, and call detail record collection. On a call-by-call basis, this traffic is not very significant. However, when a single database server is used to provide service to a large number of Gateways, this traffic will be concentrated to a single point. Therefore, if a single centralized database server will be used, it is sensible to locate it at a network hub, so that the traffic from all gateway servers will not overload a single network link.

The second kind of traffic is that between the database server and the underlying Database implementation. The database server is capable of running on a different machine than the one that runs the database. From a network standpoint, there may be motivation to do this if a database server needs to be exposed to the Internet for the purpose of CDR insertion, with the actual database (which might hold other, sensitive data) not accessible directly via the Internet. If, however, the database server and underlying database are not run on the same machine, then there will be network communication between the two machines. The volume of this traffic will depend directly on the load placed on the database server, but it will be similar in volume to the aggregated database-related traffic of all the gateway servers that the database server supports. Therefore, if possible, without violating security requirements, it is recommended that the database server be run on the same machine as the underlying database, or in a

LAN environment. Using an underlying database connected in a WAN environment with low available bandwidth may cause performance problems.

d. Management System

Using the management system to administer any of the other components of the system does not typically require a significant amount of bandwidth. The only exceptions to this are when large lists of routes are being manipulated, in which case the initial query of the routing database may require the routes to be downloaded to the management system. Although this should not normally be a problem, using a routing server that is connected to the management system either locally or via a LAN connection can alleviate the problem. Even if the routing server is operated remotely, the use of a routing server allows only one connection (the one between the routing server and management system) to become heavily utilized during the start of routing table administration.

e. Provisioning System

The provisioning system interfaces exclusively with database servers. The network requirements of this interaction depend on the types of operations being performed. Clearly, obtaining and modifying a user record do not require significant amounts of bandwidth since the amount of data contained in a single user record is relatively small. However, when doing bulk record fetching, such as fetching one hundred user records for summary display, or several hundred call records, the amount of data transferred between the database server and provisioning system may be large. If one will be using the provisioning system to interact with large amounts of data, then it is recommended that the database server be located on the same LAN (or even the same machine) as the provisioning system. If it is necessary to locate the database server and provisioning system across a WAN connection, then care should to be taken to

ensure sufficient capacity, and to prevent a conflict with time-sensitive voice data.

2. Traffic Prioritization

The network requirements of the present invention components were explained above in order to guide the formation of a network topology. However, regardless of the network topology selected, additional information can be used by capable routers in order to prioritize the different kinds of traffic. The following table shows the various port numbers that can be used to prioritize traffic in one example:

Traffic Type	Source -> Destination Port Number/Range
Gateway -> Gateway Control data	6789 -> 6789
Gateway -> Gateway Media data	User configurable range
Gateway -> Routing data	6789 -> 6791
Gateway -> Database data	6789 -> 6792
GW/RT/DB -> Management Info	6789/6791/6792 -> 6790
Routing -> Provisioning data	6791 -> xxx

From the above table, it can be observed that control oriented data originating from a particular component always originates from a single UDP/IP port. All components of a VoIP system in the present invention can have a single unique port number, as follows:

5	<u>Component</u>	<u>UDP/IP Port Number</u>
	Gateway Server	6789
	Management System	6790
	Routing Server	6791
	Database Server	6792
10	Provisioning System	xxxx

This information can be used to prioritize traffic, if the underlying routers are capable. Typically, it is desirable to ensure that Gateway to Gateway media data has the highest priority, time-sensitive control data (Gateway to Database/routing servers) has medium priority, and management/provisioning traffic has the lowest priority. This policy may of course need to be combined with the policies for other types of traffic on the same network.

V. Example Computer System

An example of a computer system 500 is shown in FIG. 5. Computer system(s) 500 can execute software to carry out any of the functionality described above with respect to VoIP system 200, including any of the servers and systems 210-240 and 280, and Management System 250, Provisioning System 260, and network manager 270.

Computer system 500 represents any single or multi-processor computer. Single-threaded and multi-threaded computers can be used. Unified or distributed memory systems can be used.

Computer system 500 includes one or more processors, such as processor 504. One or more processors 504 can execute software implementing all or part of VoIP system 200 as described above. Each processor 504 is connected to a communication infrastructure 502 (e.g., a communications bus, cross-bar, or network). Various software embodiments are described in terms of this exemplary computer system. After reading this description, it will become apparent to a person skilled in the relevant art how to implement the invention using other computer systems and/or computer architectures.

Computer system 500 also includes a main memory 508, preferably random access memory (RAM), and can also include secondary memory 510. Secondary memory 510 can include, for example, a hard disk drive 512 and/or a removable storage drive 514, representing a floppy disk drive, a magnetic tape drive, an optical disk drive, etc. The removable storage drive 514 reads from and/or writes to a removable storage unit 518 in a well known manner. Removable storage unit 518 represents a floppy disk, magnetic tape, optical disk, etc., which is read by and written to by removable storage drive 514. As will be appreciated, the removable storage unit 518 includes a computer usable storage medium having stored therein computer software and/or data.

In alternative embodiments, secondary memory 510 may include other similar means for allowing computer programs or other instructions to be loaded into computer system 500. Such means can include, for example, a removable storage unit 522 and an interface 520. Examples can include a program cartridge and cartridge interface (such as that found in video game devices), a removable memory chip (such as an EPROM, or PROM) and associated socket, and other removable storage units 522 and interfaces 520 which allow software and data to be transferred from the removable storage unit 522 to computer system 500.

Computer system 500 can also include a communications interface 524. Communications interface 524 allows software and data to be transferred between computer system 500 and external devices via communications path 526.

Examples of communications interface 524 can include a modem, a network interface (such as Ethernet card), a communications port, etc. Software and data transferred via communications interface 524 are in the form of signals which can be electronic, electromagnetic, optical or other signals capable of being received by communications interface 524, via communications path 526. Note that communications interface 524 provides a means by which computer system 500 can interface to a network such as the Internet.

The present invention can be implemented using software running (that is, executing) in an environment similar to that described above with respect to FIG. 5. In this document, the term "computer program product" is used to generally refer to removable storage unit 518, a hard disk installed in hard disk drive 512, or a carrier wave or other signal carrying software over a communication path 526 (wireless link or cable) to communication interface 524. A computer useable medium can include magnetic media, optical media, or other recordable media, or media that transmits a carrier wave. These computer program products are means for providing software to computer system 500.

Computer programs (also called computer control logic) are stored in main memory 508 and/or secondary memory 510. Computer programs can also be received via communications interface 524. Such computer programs, when executed, enable the computer system 500 to perform the features of the present invention as discussed herein. In particular, the computer programs, when executed, enable the processor 504 to perform the features of the present invention. Accordingly, such computer programs represent controllers of the computer system 500.

In an embodiment where the invention is implemented using software, the software may be stored in a computer program product and loaded into computer system 500 using removable storage drive 514, hard drive 512, or communications interface 524. Alternatively, the computer program product may be downloaded to computer system 500 over communications path 526. The

control logic (software), when executed by the one or more processors 504, causes the processor(s) 504 to perform the functions of the invention as described herein.

In another embodiment, the invention is implemented primarily in firmware and/or hardware using, for example, hardware components such as application specific integrated circuits (ASICs). Implementation of a hardware state machine so as to perform the functions described herein will be apparent to persons skilled in the relevant art(s).

VI. Conclusion

While specific embodiments of the present invention have been described above, it should be understood that they have been presented by way of example only, and not limitation. It will be understood by those skilled in the art that various changes in form and details may be made therein without departing from the spirit and scope of the invention as defined in the appended claims. Thus, the breadth and scope of the present invention should not be limited by any of the above-described exemplary embodiments, but should be defined only in accordance with the following claims and their equivalents.

What Is Claimed Is:

1 1. A system for providing database services in a voice
2 communication system over a packet-switched network, comprising:
3 at least one database that stores data;
4 at least one database client server that accesses said at least one database;
5 at least one database client that accesses said at least one database client
6 server;
7 at least one service module that accesses the database system;
8 a multi-database client that manages said at least one service module's
9 access to said at least one database client; and
10 a database access manager that monitors said at least one service module
11 access to said multi-database client.

1 2. A system of claim 1, wherein said at least one service module
2 includes an account service.

1 3. A system of claim 1, wherein said at least one service module
2 includes an call detail record service.

1 4. A system of claim 1, wherein said at least one service module
2 includes an event log service.

1 5. A system of claim 1, wherein said at least one database includes
2 a relational database.

1 6. A system of claim 1, wherein said at least one database includes
2 a non-relational database.

1 7. A system of claim 1, wherein said at least one database and said
2 at least one database client server are configured behind a firewall.

1 8. A system of claim 7, wherein said at least one database client
2 accesses said at least one database client server through said firewall.

1 9. A system of claim 2, wherein said service module for managing
2 accounts comprises:

3 obtaining means for obtaining account information;

4 requesting means for requesting said account information from at least
5 one multi-database client;

6 receiving means for receiving said account information from at least one
7 multi-database client;

8 first forwarding means for forwarding said account information to a multi-
9 database client; and

10 second forwarding means for forwarding said account information to a
11 database access manager.

1 10. A system of claim 9, wherein said account information includes
2 a user's name, telephone number, identification number, password, system
3 service access level, billing information, and preferred calling patterns.

1 11. A system of claim 9, wherein said account information includes
2 a group's name, identification number, password, system service access level,
3 usage restrictions, billing information, and preferred calling patterns.

1 12. A system of claim 3, wherein said service module for recording
2 call details comprises:

3 recording means for recording call details;
4 storing means for storing said call details in a call detail records server;
5 first sending means for sending said call details to a multi-client database;
6 receiving means for receiving a request for said stored call detail records
7 from a database access manager; and
8 second sending means for sending said stored called details to said multi-
9 database client from said call detail records server.

1 13. A system of claim 12, wherein said storing means comprises:
2 a local cache to automatically store call detail records,
3 wherein said service module for recording call details stores said call
4 details when first sending means is unsuccessful.

1 14. A system of claim 4, wherein said service module for logging
2 events comprises:
3 receiving means for receiving a request for service;
4 determining means for determining the type of service being requested;
5 recording means for recording said request for service and the details of
6 the service provided as an event; and
7 sending means for sending said event to a multi-database client.

1 15. A system of claim 1, wherein said database access manager
2 comprises:
3 connecting means for connecting service modules to a multi-database
4 client;
5 monitoring means for monitoring requests for service from said service
6 modules to a multi-database client;
7 authorizing means for authorizing said requests for service from said
8 service modules to said multi-database client; and

9 prioritizing means for prioritizing said requests for service from said
10 service modules to said multi-database client.

1 16. A system of claim 1, wherein said multi-database client
2 comprises:
3 first connecting means for connecting with service modules;
4 second connecting means for connecting with at least one database client;
5 receiving means for receiving requests for service from said service
6 modules;
7 queuing means for queuing said requests for service from said service
8 modules;
9 forwarding means for forwarding requests for service from said service
10 modules to said at least one database client; and
11 confirming means for confirming with said at least one database client
12 that said requests for service from said service modules was received as
13 forwarded.

1 17. A system of claim 1, further comprising:
2 at least one database client server that receives requests from said multi-
3 database client;
4 at least one database client that receives said requests from said at least
5 one database client server; and
6 at least one database that receives said requests from said at least one
7 database client.

1 18. A system of claim 17, wherein said at least one database client
2 server comprises:
3 a software environment that provides a common platform for services;
4 a database manager that communicates with at least one database client;

5 a database service that instantiates said database manager,
6 a thread pool queue that receives requests from said database manager;
7 a request handler that receives said requests from said thread pool queue;
8 an internal services manager that manages said request handler's access
9 to service modules;
10 a database connection that links said service modules with a database
11 interface; and
12 a database connection manager that manages said database connection.

1 19. A system of claim 17, wherein said at least one database client
2 servers comprises:

3 executing means for executing a database service;
4 instantiating means for instantiating a database manager to receive and
5 respond to requests;
6 building means for building a thread pool queue to queue said requests;
7 first connecting means for connecting to a software environment;
8 second connecting means for connecting to a database client;
9 maintaining means for maintaining said thread pool queue;
10 first receiving means for receiving said requests from a request handler;
11 second receiving means for receiving said requests from said database
12 client;
13 responding means for responding to said database client;
14 first forwarding means for forwarding said requests to said thread pool
15 queue;
16 second forwarding means for forwarding said requests to said software
17 environment; and
18 third forwarding means for forwarding said requests to a multi-database
19 client.

1 20. A system of claim 18, wherein said internal service manager
2 comprises:

3 first connecting means for connecting to a requests handler;
4 second connecting means for connecting to at least one service module;
5 third connecting means for connecting to a multi-database client;
6 fourth connecting means for connecting to a database access manager,
7 forwarding means for forwarding requests for service from said at least
8 one service module to said requests handler; and
9 confirming means for confirming the handling of said requests for service
10 by said service modules.

1 21. A system of claim 17, wherein said at least one database client
2 comprises:

3 first connecting means for connecting to a database client server;
4 second connecting means for connecting to a multi-database client,
5 third connecting means for connecting to a software environment;
6 first receiving means for receiving requests for service from said multi-
7 database client;
8 second receiving means for receiving said requests for service from said
9 software environment;
10 first forwarding means for forwarding said requests for service to said
11 database client server;
12 third receiving means for receiving a response from said database client
13 server;
14 second forwarding means for forwarding said response to said multi-
15 database client; and
16 third forwarding means for forwarding said response to said software
17 environment.

1 22. A system for provisioning, comprising:

2 first instantiating means for instantiating a provisioning system document
3 that provides query objects with respect to record in a database;

4 first connecting means for connecting said provisioning system document
5 to a software environment;

6 second connecting means for connecting a provisioning system document
7 to a multi-database client;

8 second instantiating means for instantiating a provisioning system view
9 that provides an interface with respect to said provisioning system document;

10 third connecting means for connecting said provisioning system view with
11 said provisioning system document;

12 providing means for providing information from a database client
13 subsystem where said database client subsystem is connected to said database.

1 23. A system for management, comprising:

2 a software environment that provides a communications layer for a
3 plurality of components of a voice over internet communications system,

4 a management system document that provides configuration objects with
5 respect to said plurality of components of a voice over internet communications
6 system;

7 a management system view that provides an interface with respect to said
8 management system document;

9 a management client that interacts with said plurality of components of
10 a voice over internet communications system;

11 management client objects that enable management events;

12 management controls that interface with said management client objects,
13 said management controls monitoring said management events,

14 wherein said management events are sent between said management client
15 and said management client objects and said management client objects and said

16 management controls with respect to interaction between said management client
17 and said management system view.

1 24. A system of claim 23, further comprising:
2 a cache of object controls,
3 wherein said cache of object controls receives updates from said
4 management client.

1 25. A system of claim 24, wherein said cache of object controls
2 implements updates with respect to said management client objects and said
3 management controls.

1 26. A system for managing a network of components in a software
2 environment, comprising:
3 querying means for querying a client database for configuration
4 information;
5 receiving means for receiving data from said client database;
6 determining means for determining if said data represents updated
7 configuration information; and
8 sending means for sending a message representative of said updated
9 configuration information to said software environment.

1 27. A system of claim 26, wherein said software environment includes
2 a gateway server.

1 28. A system of claim 26, wherein said software environment includes
2 a database server.

1 29. A system of claim 26, wherein said software environment includes
2 a routing server.

1 30. A method of providing database services in a voice
2 communication system over a packet-switched network, comprising the steps of:
3 storing data in at least one database;
4 accessing said at least one database with at least one database client
5 server;
6 accessing said at least one database client server with at least one database
7 client;
8 accessing the database system with at least one service module;
9 managing said at least one service module's access to said at least one
10 database client with a multi-database client; and
11 monitoring said at least one service module access to said multi-database
12 client a database access manager.

1 31. A method of claim 30, wherein said at least one service module
2 includes an account service.

1 32. A method of claim 30, wherein said at least one service module
2 includes an call detail record service.

1 33. A method of claim 30, wherein said at least one service module
2 includes an event log service.

1 34. A method of claim 30, wherein said at least one database includes
2 a relational database.

1 35. A method of claim 30, wherein said at least one database includes
2 a non-relational database.

1 36. A method of claim 30, wherein said at least one database and said
2 at least one database client server are configured behind a firewall.

1 37. A method of claim 36, wherein said at least one database client
2 accesses said at least one database client server through said firewall.

1 38. A method of claim 31, wherein said service module for managing
2 accounts comprises the steps of:
3 obtaining account information;
4 requesting said account information from at least one multi-database
5 client;
6 receiving said account information from at least one multi-database client;
7 for forwarding said account information to a multi-database client, and
8 for forwarding said account information to a database access manager.

1 39. A method of claim 38, wherein said account information includes
2 a user's name, telephone number, identification number, password, system service
3 access level, billing information, and preferred calling patterns.

1 40. A method of claim 38, wherein said account information includes
2 a group's name, identification number, password, system service access level,
3 usage restrictions, billing information, and preferred calling patterns.

1 41. A method of claim 32, wherein said service module for recording
2 call details comprises the steps of:
3 recording call details;

4 storing said call details in a call detail records server;
5 sending said call details to a multi-client database,
6 receiving a request for said stored call detail records from a database
7 access manager; and
8 sending said stored called details to said multi-database client from said
9 call detail records server.

1 42. A method of claim 41, wherein said storing means comprises the
2 steps of:
3 caching locally and automatically call detail records,
4 wherein said service module for recording call details stores said call
5 details when first sending means is unsuccessful.

1 43. A method of claim 33, wherein said service module for logging
2 events comprises the steps of:
3 receiving a request for service;
4 determining the type of service being requested;
5 recording said request for service and the details of the service provided
6 as an event; and
7 sending said event to a multi-database client.

1 44. A method of claim 30, wherein said database access manager
2 comprises the steps of:
3 connecting service modules to a multi-database client;
4 monitoring requests for service from said service modules to a
5 multi-database client;
6 authorizing said requests for service from said service modules to said
7 multi-database client; and

8 prioritizing said requests for service from said service modules to said
9 multi-database client.

1 45. A method of claim 30, wherein said multi-database client
2 comprises the steps of:

3 connecting with service modules;
4 connecting with at least one database client;
5 receiving requests for service from said service modules;
6 queuing said requests for service from said service modules;
7 forwarding requests for service from said service modules to said at least
8 one database client; and

9 confirming with said at least one database client that said requests for
10 service from said service modules was received as forwarded.

1 46. A method of claim 30, further comprising the steps of
2 receiving requests from said multi-database client with at least one
3 database client server;

4 receiving said requests from said at least one database client server with
5 at least one database client; and

6 receiving said requests from said at least one database client with at least
7 one database.

1 47. A method of claim 46, wherein said at least one database client
2 server comprises the steps of:

3 providing a common platform for services with a software environment;
4 communicating with at least one database client with a database manager;
5 instantiating said database manager with a database service;
6 receiving requests from said database manager with a thread pool queue;

7 receiving said requests from said thread pool queue with a request
8 handler;
9 managing said request handler's access to service modules with an internal
10 services manager;
11 linking said service modules with a database interface with a database
12 connection; and
13 managing said database connection with a database connection manager.

1 48. A method of claim 46, wherein said at least one database client
2 servers comprises the steps of:

3 executing a database service;
4 instantiating a database manager to receive and respond to requests;
5 building a thread pool queue to queue said requests;
6 connecting to a software environment;
7 connecting to a database client;
8 maintaining said thread pool queue;
9 receiving said requests from a request handler;
10 receiving said requests from said database client;
11 responding to said database client;
12 forwarding said requests to said thread pool queue;
13 forwarding said requests to said software environment; and
14 forwarding said requests to a multi-database client.

1 49 A method of claim 47, wherein said internal service manager
2 comprises the steps of:

3 connecting to a requests handler;
4 connecting to at least one service module;
5 connecting to a multi-database client;
6 connecting to a database access manager;

7 forwarding requests for service from said at least one service module to
8 said requests handler; and
9 confirming the handling of said requests for service by said service
10 modules.

1 50. A method of claim 46, wherein said at least one database client
2 comprises the steps of:
3 connecting to a database client server;
4 connecting to a multi-database client;
5 connecting to a software environment;
6 receiving requests for service from said multi-database client;
7 receiving said requests for service from said software environment;
8 forwarding said requests for service to said database client server,
9 receiving a response from said database client server;
10 forwarding said response to said multi-database client; and
11 forwarding said response to said software environment.

1 51. A method of provisioning, comprising the steps of:
2 instantiating a provisioning system document that provides query objects
3 with respect to record in a database;
4 connecting said provisioning system document to a software environment;
5 connecting a provisioning system document to a multi-database client;
6 instantiating a provisioning system view that provides an interface with
7 respect to said provisioning system document;
8 connecting said provisioning system view with said provisioning system
9 document;
10 providing information from a database client subsystem where said
11 database client subsystem is connected to said database.

1 52. A method of management, comprising the steps of:
2 providing a communications layer for a plurality of components of a voice
3 over internet communications system with a software environment;
4 providing configuration objects with respect to said plurality of
5 components of a voice over internet communications system with a management
6 system document,
7 providing an interface with respect to said management system document
8 with a management system view;
9 interacting with said plurality of components of a voice over internet
10 communications system with a management client;
11 enabling management events with management client objects;
12 interfacing with said management client objects, said management
13 controls monitoring said management events with management controls,
14 wherein said management events are sent between said management client
15 and said management client objects and said management client objects and said
16 management controls with respect to interaction between said management client
17 and said management system view.

1 53. A method of claim 52, further comprising the steps of:
2 caching object controls, wherein said object controls are updated from
3 said management client.

1 54. A method of claim 53, wherein said object controls are
2 implemented with respect to said management client objects and said
3 management controls when updated.

1 55. A method of managing a network of components in a software
2 environment, comprising the steps of:
3 querying a client database for configuration information;

4 receiving data from said client database;
5 determining if said data represents updated configuration information; and
6 sending a message representative of said updated configuration
7 information to said software environment.

1 56. A method of claim 55, wherein said software environment
2 includes a gateway server.

1 57. A method of claim 55, wherein said software environment
2 includes a database server.

1 58. A method of claim 55, wherein said software environment
2 includes a routing server.

1 59. A computer program product comprising a computer useable
2 medium having computer program logic stored therein, said computer program
3 logic comprising:

4 means for enabling a computer to access at least one database that stores
5 data;

6 means for enabling a computer to access at least one database client server
7 that accesses said at least one database;

8 means for enabling a computer to access at least one database client that
9 accesses said at least one database client server;

10 means for enabling a computer to access at least one service module that
11 accesses the database system;

12 means for enabling a computer to access a multi-database client that
13 manages said at least one service module's access to said at least one database
14 client; and

15 means for enabling a computer to access a database access manager that
16 monitors said at least one service module access to said multi-database client.

1 60. The computer program product of claim 59, wherein said at least
2 one service module includes an account service.

1 61. The computer program product of claim 59, wherein said at least
2 one service module includes an call detail record service.

1 62. The computer program product of claim 59, wherein said at least
2 one service module includes an event log service.

1 63. The computer program product of claim 59, wherein said at least
2 one database includes a relational database.

1 64. The computer program product of claim 59, wherein said at least
2 one database includes a non-relational database.

1 65. The computer program product of claim 59, wherein said at least
2 one database and said at least one database client server are configured behind a
3 firewall

1 66. The computer program product of claim 65, wherein said at least
2 one database client accesses said at least one database client server through said
3 firewall.

1 67. The computer program product of claim 60, wherein said service
2 module for managing accounts comprises:

3 means for enabling the computer to obtain account information;

4 means for enabling the computer to request said account information from
5 at least one multi-database client;

6 means for enabling the computer to receive said account information from
7 at least one multi-database client;

8 means for enabling the computer to forward said account information to
9 a multi-database client; and

10 means for enabling the computer to forward said account information to
11 a database access manager.

1 68. The computer program product of claim 67, wherein said account
2 information includes a user's name, telephone number, identification number,
3 password, system service access level, billing information, and preferred calling
4 patterns.

1 69. The computer program product of claim 67, wherein said account
2 information includes a group's name, identification number, password, system
3 service access level, usage restrictions, billing information, and preferred calling
4 patterns.

1 70. The computer program product of claim 61 wherein said service
2 module for recording call details comprises:

3 means for enabling the computer to record call details;

4 means for enabling the computer to store said call details in a call detail
5 records server;

6 means for enabling the computer to send said call details to a multi-client
7 database;

8 means for enabling the computer to receive a request for said stored call
9 detail records from a database access manager; and

10 means for enabling the computer to send said stored called details to said
11 multi-database client from said call detail records server.

1 71. The computer program product of claim 70, wherein said storing
2 means comprises:

3 means for enabling the computer to locally and automatically store call
4 detail records in a cache,

5 wherein said service module for recording call details stores said call
6 details when first sending means is unsuccessful.

1 72. The computer program product of claim 62, wherein said service
2 module for logging events comprises:

3 means for enabling the computer to receive a request for service;

4 means for enabling the computer to determine the type of service being
5 requested;

6 means for enabling the computer to record said request for service and the
7 details of the service provided as an event; and

8 means for enabling the computer to send said event to a multi-database
9 client.

1 73. The computer program product of claim 59, wherein said database
2 access manager comprises:

3 means for enabling the computer to connect service modules to a
4 multi-database client;

5 means for enabling the computer to monitor requests for service from said
6 service modules to a multi-database client;

7 means for enabling the computer to authorize said requests for service
8 from said service modules to said multi-database client; and

9 means for enabling the computer to prioritize said requests for service
10 from said service modules to said multi-database client.

1 74. The computer program product of claim 59, wherein said
2 multi-database client comprises:

3 means for enabling the computer to connect with service modules;

4 means for enabling the computer to connect with at least one database
5 client;

6 means for enabling the computer to receive requests for service from said
7 service modules;

8 means for enabling the computer to queue said requests for service from
9 said service modules;

10 means for enabling the computer to forward requests for service from said
11 service modules to said at least one database client; and

12 means for enabling the computer to confirm with said at least one
13 database client that said requests for service from said service modules was
14 received as forwarded.

1 75. The computer program product of claim 59, further comprising:

2 means for enabling a computer to access at least one database client server
3 that receives requests from said multi-database client;

4 means for enabling a computer to access at least one database client that
5 receives said requests from said at least one database client server; and

6 means for enabling a computer to access at least one database that
7 receives said requests from said at least one database client.

1 76. The computer program product of claim 75, wherein said at least
2 one database client server comprises:

3 means for enabling a computer to access a software environment that
4 provides a common platform for services;

5 means for enabling a computer to access a database manager that
6 communicates with at least one database client;

7 means for enabling a computer to access a database service that
8 instantiates said database manager;

9 means for enabling a computer to access a thread pool queue that receives
10 requests from said database manager;

11 means for enabling a computer to access a request handler that receives
12 said requests from said thread pool queue;

13 means for enabling a computer to access an internal services manager that
14 manages said request handler's access to service modules;

15 means for enabling a computer to access a database connection that links
16 said service modules with a database interface; and

17 means for enabling a computer to access a database connection manager
18 that manages said database connection.

1 77. The computer program product of claim 75, wherein said at least
2 one database client servers comprises:

3 means for enabling the computer to execute a database service;

4 means for enabling the computer to instantiate a database manager to
5 receive and respond to requests;

6 means for enabling the computer to build a thread pool queue to queue
7 said requests;

8 means for enabling the computer to connect to a software environment;

9 means for enabling the computer to connect to a database client;

10 means for enabling the computer to maintain said thread pool queue;

11 means for enabling the computer to receive said requests from a request
12 handler;

13 means for enabling the computer to receive said requests from said
14 database client;
15 means for enabling the computer to respond to said database client;
16 means for enabling the computer to forward said requests to said thread
17 pool queue;
18 means for enabling the computer to forward said requests to said software
19 environment; and
20 means for enabling the computer to forward said requests to a
21 multi-database client.

1 78. The computer program product of claim 76, wherein said internal
2 service manager comprises:
3 means for enabling the computer to connect to a requests handler;
4 means for enabling the computer to connect to at least one service
5 module;
6 means for enabling the computer to connect to a multi-database client;
7 means for enabling the computer to connect to a database access manager;
8 means for enabling the computer to forward requests for service from said
9 at least one service module to said requests handler; and
10 means for enabling the computer to confirm the handling of said requests
11 for service by said service modules.

1 79. The computer program product of claim 75, wherein said at least
2 one database client comprises:
3 means for enabling the computer to connect to a database client server;
4 means for enabling the computer to connect to a multi-database client;
5 means for enabling the computer to connect to a software environment;
6 means for enabling the computer to receive requests for service from said
7 multi-database client;

8 means for enabling the computer to receive said requests for service from
9 said software environment;
10 means for enabling the computer to forward said requests for service to
11 said database client server;
12 means for enabling the computer to receive a response from said database
13 client server;
14 means for enabling the computer to forward said response to said
15 multi-database client; and
16 means for enabling the computer to forward said response to said software
17 environment.

1 80. A system for provisioning, comprising:

2 means for enabling the computer to instantiate a provisioning system
3 document that provides query objects with respect to record in a database;

4 means for enabling the computer to connect said provisioning system
5 document to a software environment;

6 means for enabling the computer to connect a provisioning system
7 document to a multi-database client;

8 means for enabling the computer to instantiate a provisioning system view
9 that provides an interface with respect to said provisioning system document;

10 means for enabling the computer to connect said provisioning system
11 view with said provisioning system document;

12 means for enabling the computer to provide information from a database
13 client subsystem where said database client subsystem is connected to said
14 database.

1 81. A system for management, comprising:

2 means for enabling a computer to access a software environment that
3 provides a communications layer for a plurality of components of a voice over
4 internet communications system;

5 means for enabling a computer to access a management system document
6 that provides configuration objects with respect to said plurality of components
7 of a voice over internet communications system;

8 means for enabling a computer to access a management system view that
9 provides an interface with respect to said management system document;

10 means for enabling a computer to access a management client that
11 interacts with said plurality of components of a voice over internet
12 communications system;

13 means for enabling a computer to access management client objects that
14 enable management events;

15 means for enabling a computer to access management controls that
16 interface with said management client objects, said management controls
17 monitoring said management events,

18 wherein said management events are sent between said management client
19 and said management client objects and said management client objects and said
20 management controls with respect to interaction between said management client
21 and said management system view.

1 82. The computer program product of claim 81, further comprising.
2 means for enabling a computer to access a cache of object controls,
3 wherein said cache of object controls receives updates from said
4 management client.

1 83. The computer program product of claim 82, wherein said cache
2 of object controls implements updates with respect to said management client
3 objects and said management controls.

1 84. A system for managing a network of components in a software
2 environment, comprising:

3 means for enabling the computer to query a client database for
4 configuration information;

5 means for enabling the computer to receive data from said client database;

6 means for enabling the computer to determine if said data represents
7 updated configuration information; and

8 means for enabling the computer to send a message representative of said
9 updated configuration information to said software environment.

1 85. The computer program product of claim 84, wherein said software
2 environment includes a gateway server.

1 86. The computer program product of claim 84, wherein said software
2 environment includes a database server.

1 87. The computer program product of claim 84, wherein said software
2 environment includes a routing server.

1 88. The database services according to claim 1, further including
2 computer interface means for displaying said database services to a
3 predetermined computer, information being exchanged with said predetermined
4 computer, said computer interface means being capable of configuring said
5 database services.

1 89. The provisioning system according to claim 22, further including
2 computer interface means for displaying said provisioning system to a
3 predetermined computer, information being exchanged with said predetermined

4 computer, said computer interface means being capable of configuring said
5 provisioning system.

1 90. The management system according to claim 23, further including
2 computer interface means for displaying said management system to a
3 predetermined computer, information being exchanged with said predetermined
4 computer, said computer interface means being capable of configuring said
5 management system.

1 91. The network manager according to claim 26, further including
2 computer interface means for displaying said network manager to a
3 predetermined computer, information being exchanged with said predetermined
4 computer, said computer interface means being capable of configuring said
5 network manager.

Method, System, and Computer Program Product for Managing Database Servers and Services

Abstract

A system, method, and computer program product for managing database
5 servers and services is disclosed, which expands the capabilities of existing
systems by providing better scalability and security. A database management
system includes databases consisting of database clients and servers that handle
the storing of information, a multi-database client that manages database servers,
10 service modules which allow other systems to access the database servers, and a
database access manager that manages requests from the service modules. A
database access manager facilitates and coordinates interaction between the
service modules and the multi-database client. Service modules provide
interfaces from which other external or internal systems can access the managed
15 databases. Requests can be sent between systems and the database access
manager via the service modules. Two present and representative
implementations of service modules are: user management and call detail record
(CDR) collection. The user management service module allows systems to
access user account information stored in the databases. The CDR collection
service module inserts details for each call made on the network.

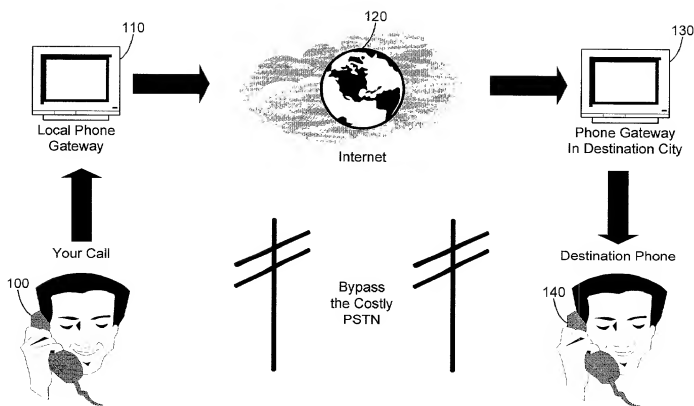


FIG. 1A




FIG. 1B

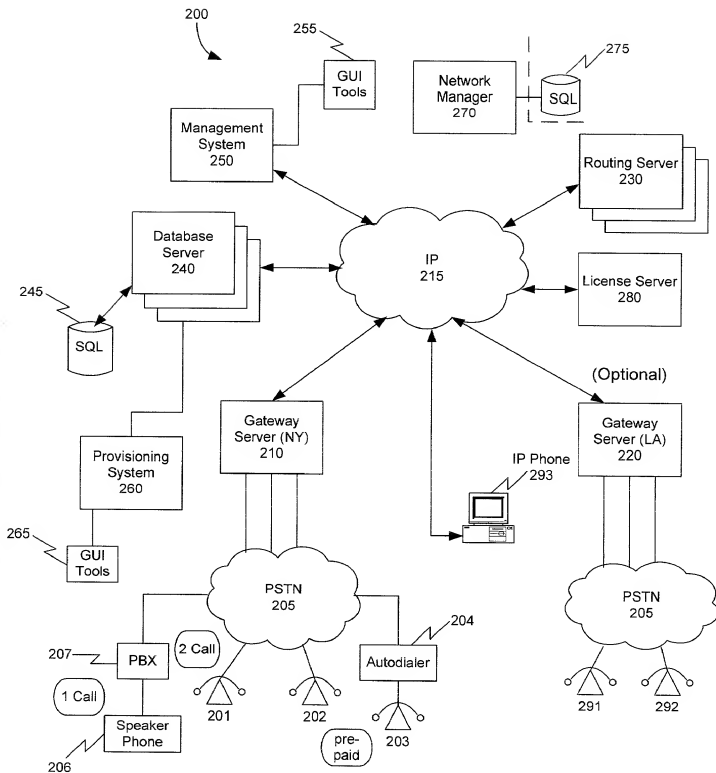


FIG. 2

210 →




FIG. 3A

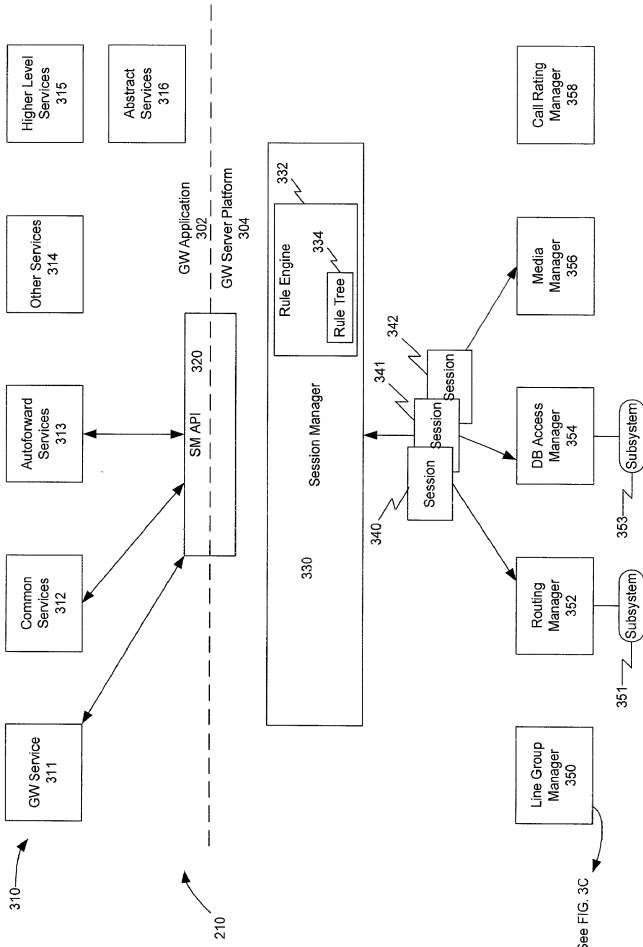


FIG. 3B

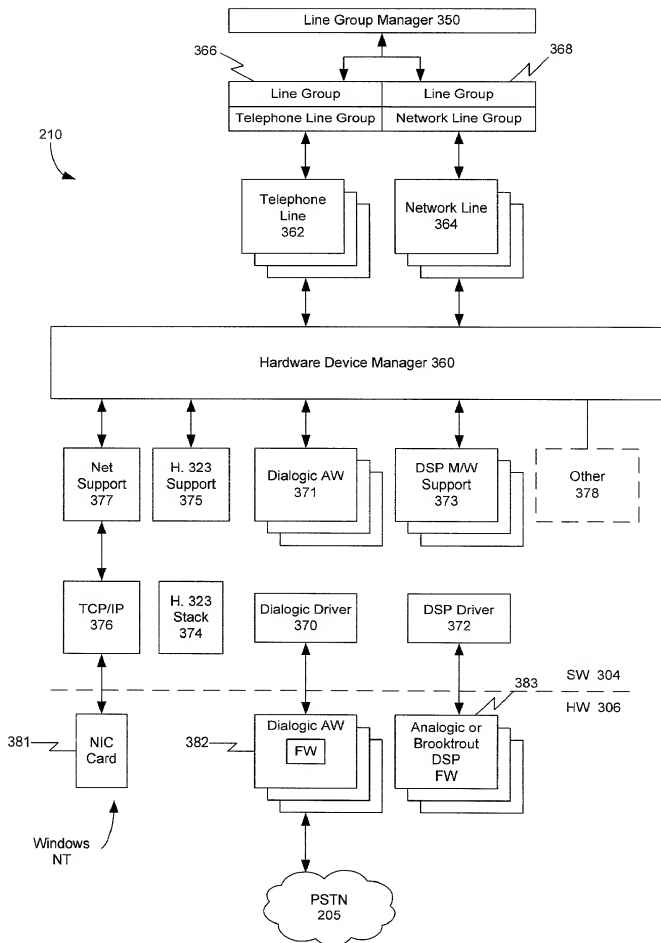


FIG. 3C

230 →




FIG. 3D

240

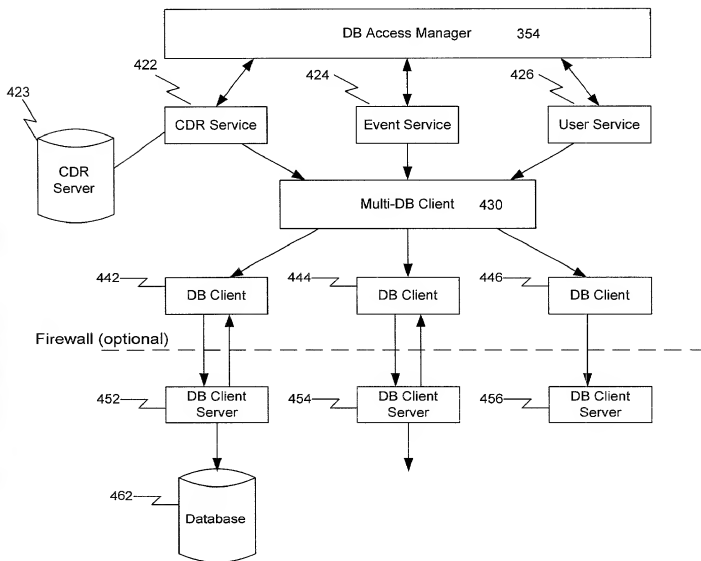


FIG. 4

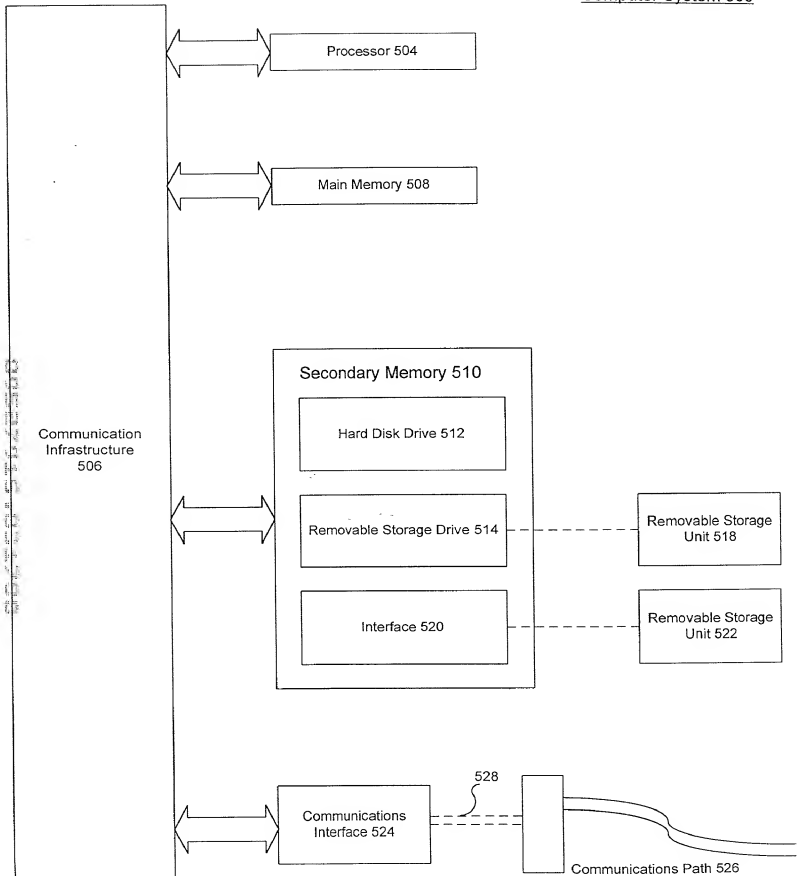


FIG. 5

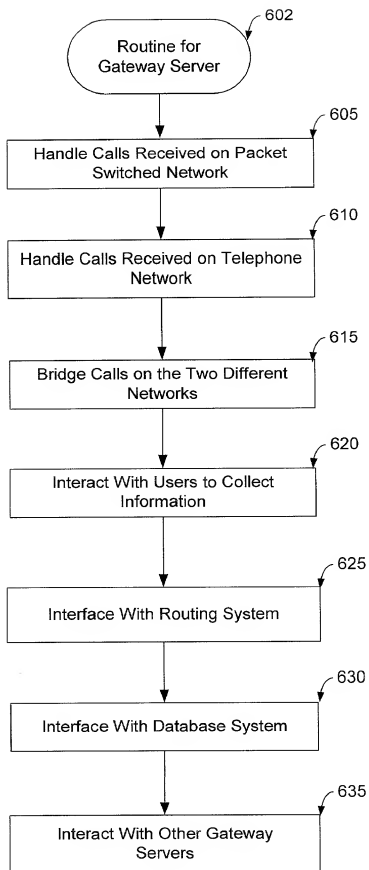


FIG. 6A

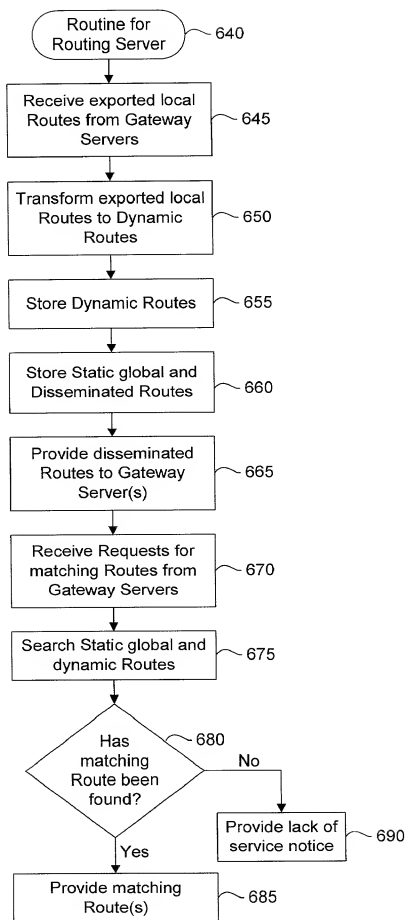


FIG. 6B

645

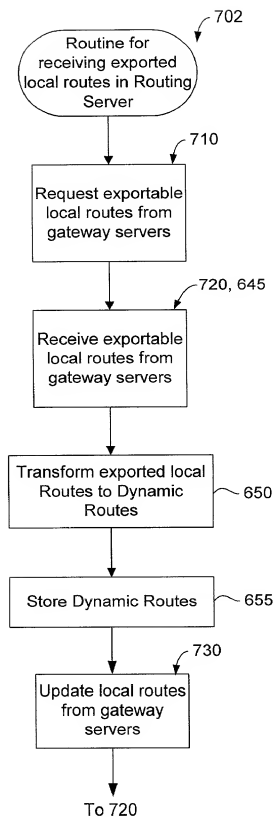


FIG. 7

650

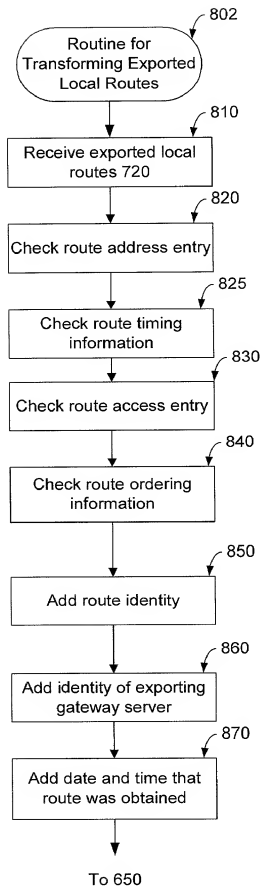


FIG. 8

665




FIG. 9

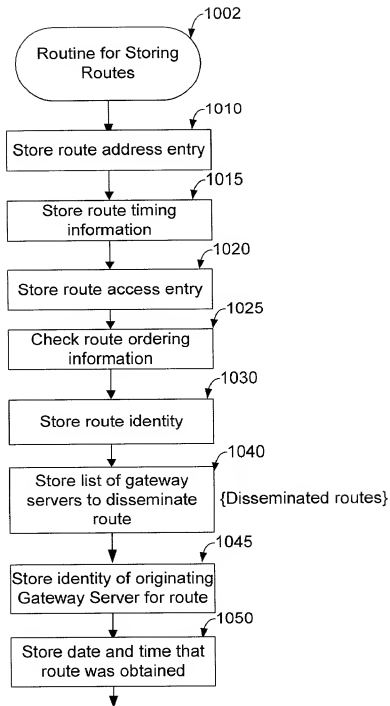


FIG. 10

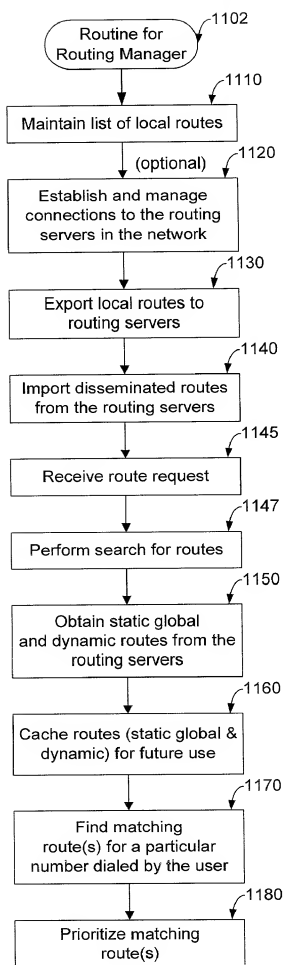
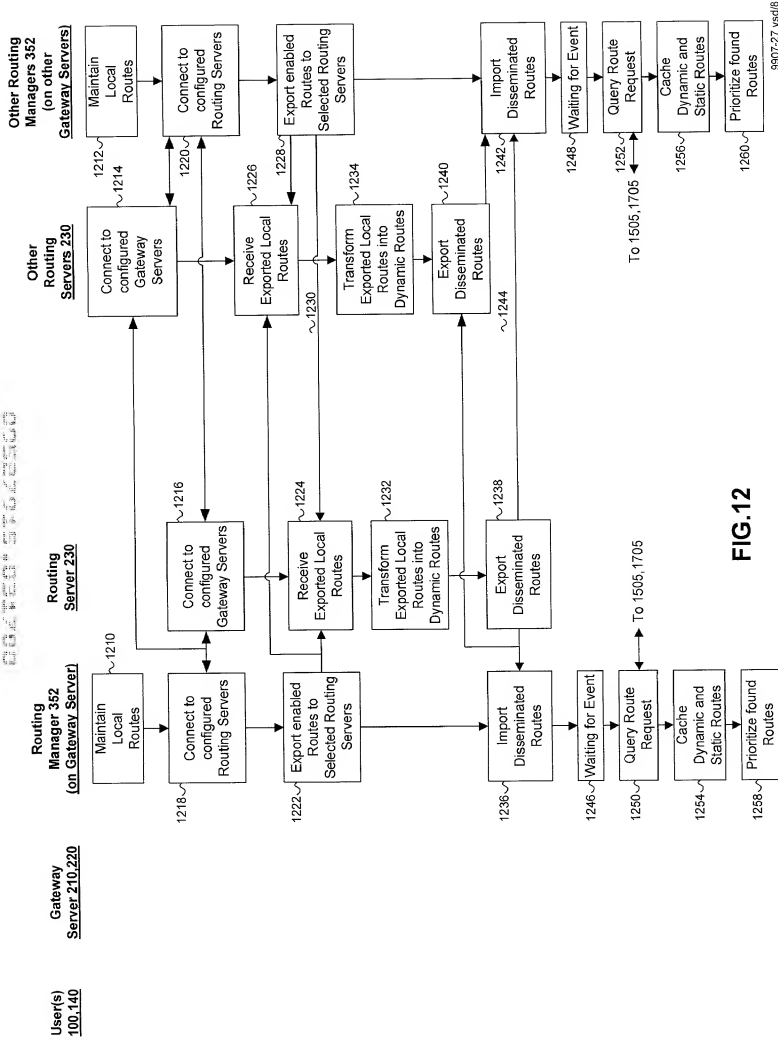


FIG. 11



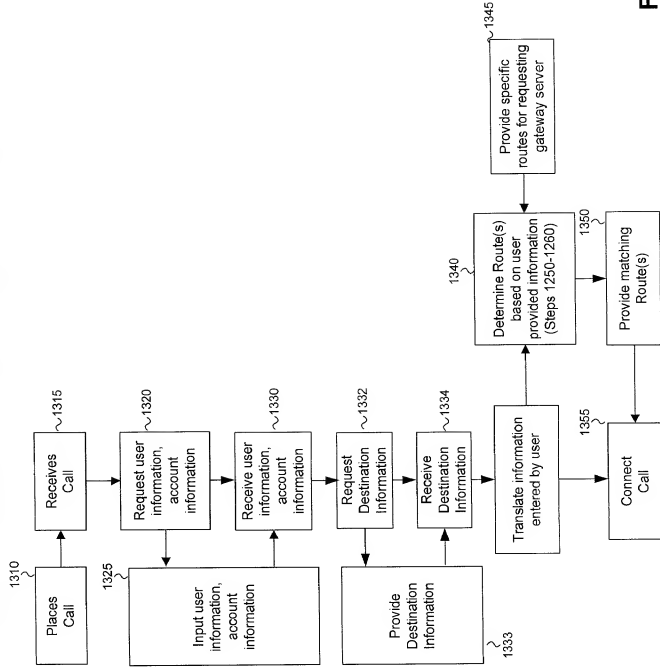


FIG.13

Other Routing
Managers 352
(on other
Gateway Servers)

Other
Routing
Servers 230

Routing
Manager 352
(on Gateway Server)

User(s)
100, 140

Gateway
Server 210, 220

Routing
Server 230

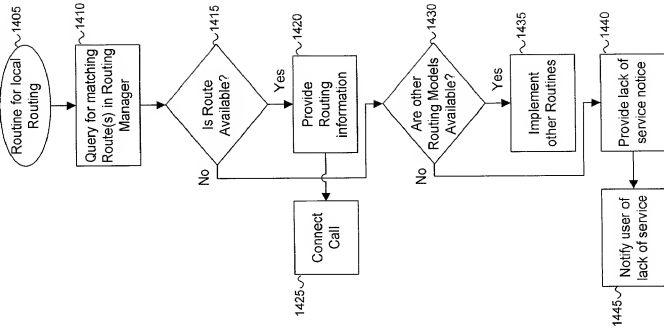


FIG. 14

Other Routing
Managers 352
(on other
Gateway Servers)

Other
Routing
Servers 230

Routing
Server 230

Routing
Manager 352
(on Gateway Server)

Gateway
Server 210, 220

User(s)
100, 140

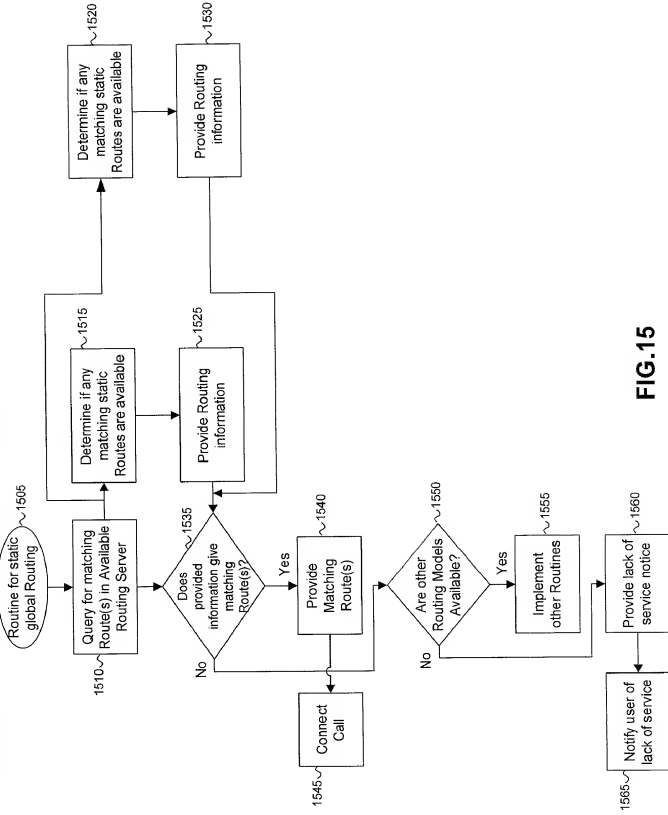


FIG. 15

User(s)
 100,140
 Gateway
 Server 210,220
 Routing
 Manager 352
 (on Gateway Server)
 Other
 Routing
 Servers 230
 (on other
 Gateway Servers)
 Other Routing
 Managers 352
 (on other
 Gateway Servers)

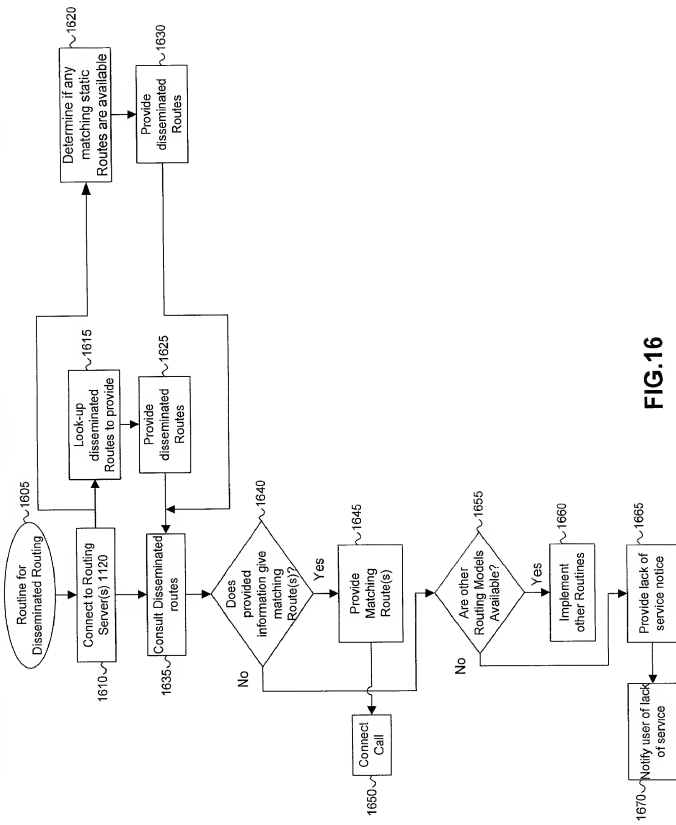


FIG. 16

Other Routing
Managers 352
(on other
Gateway Servers)

Other
Routing
Servers 230

Routing
Server 230

Routing
Manager 352
(on Gateway Server)

Gateway
Server 210,220

User(s)
100,140

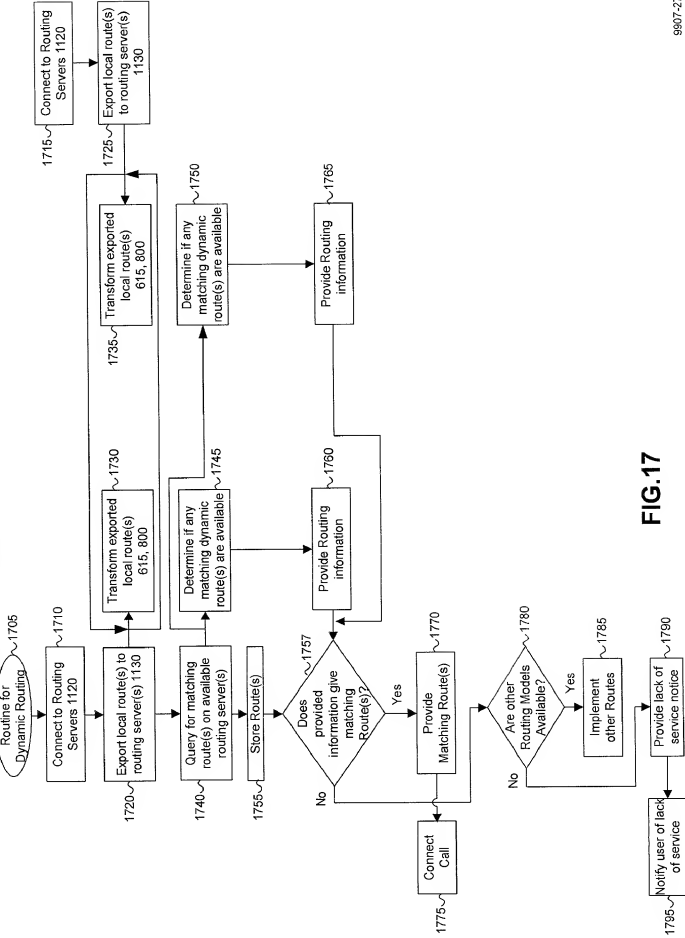


FIG.17

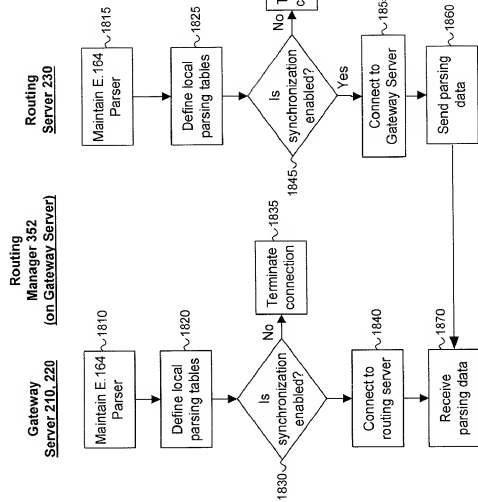


FIG. 18

1170

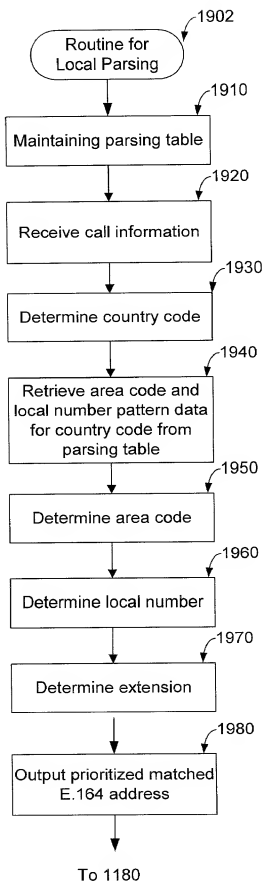
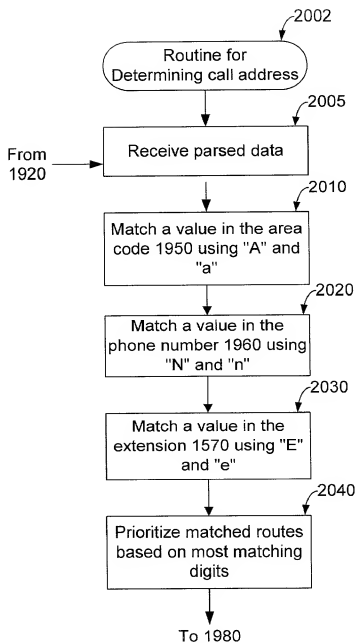


FIG. 19



Wildcard Embodiment

FIG. 20

2040

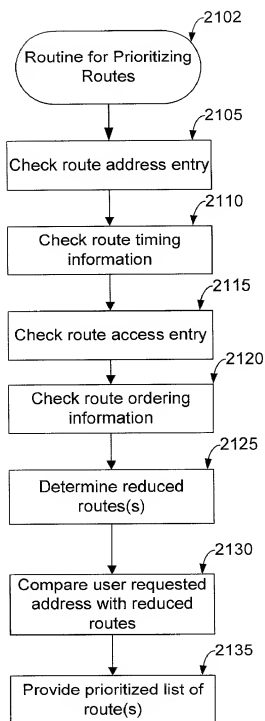


FIG. 21A

2120

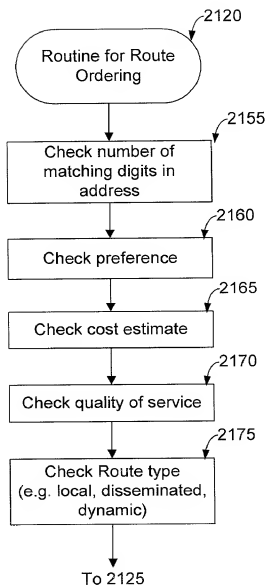


FIG. 21B




Fig. 22A




Fig. 22B




Fig. 22C

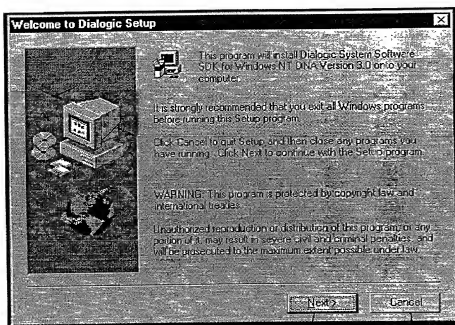


Fig. 23




Fig. 24

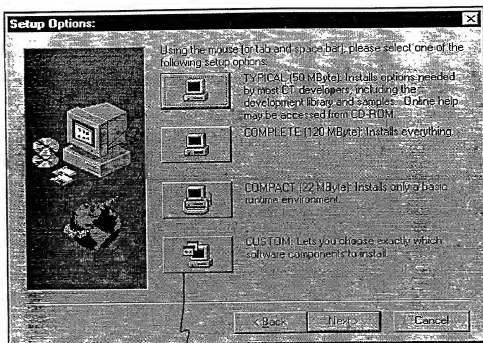


Fig. 25

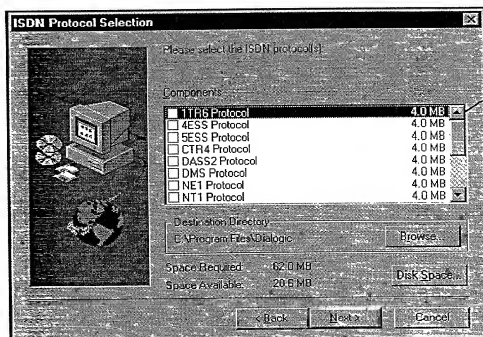


Fig. 26

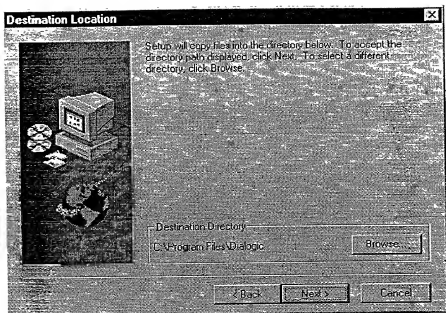


Fig. 27

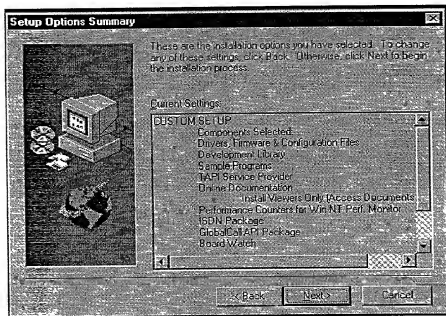


Fig. 28

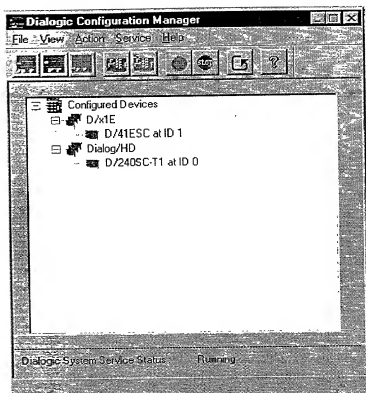


Fig. 29

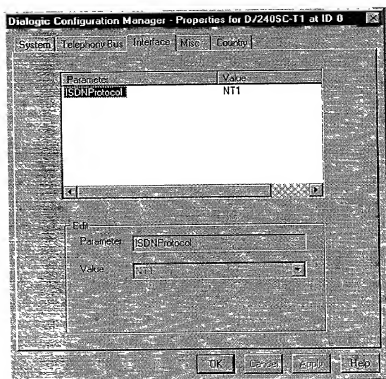


Fig. 30

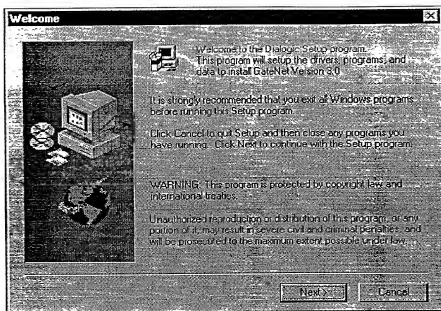


Fig.31

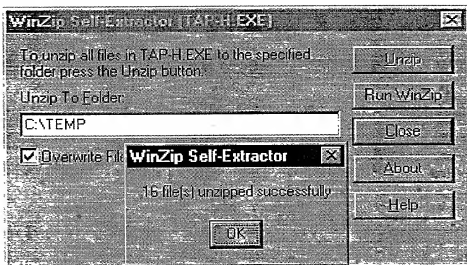


Fig. 32

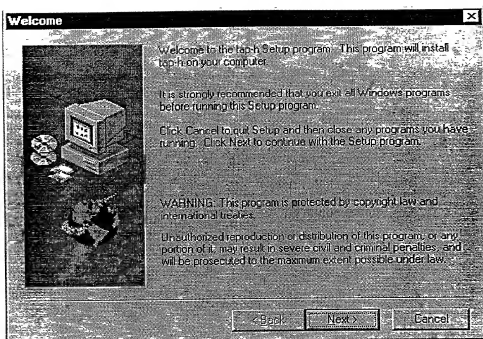


Fig. 33

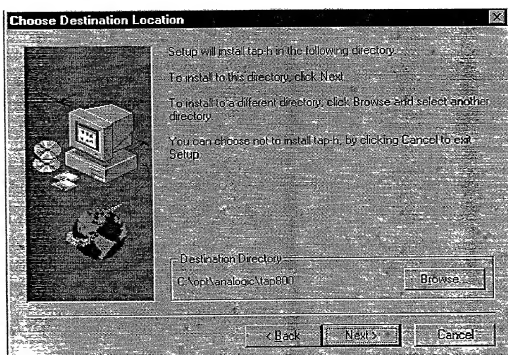
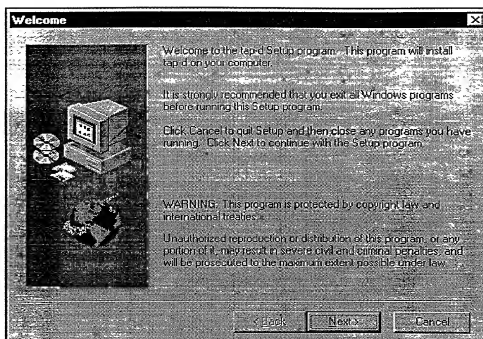


Fig. 34



~3502

Fig. 35

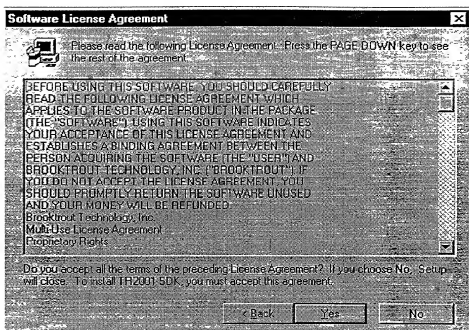


Fig. 36

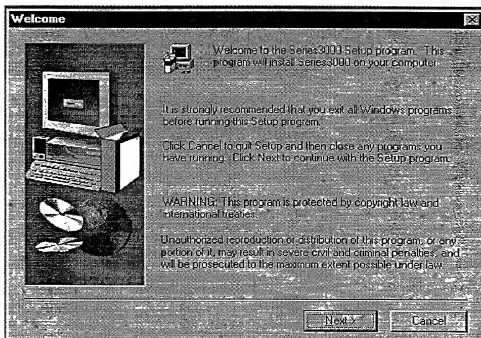


Fig. 37

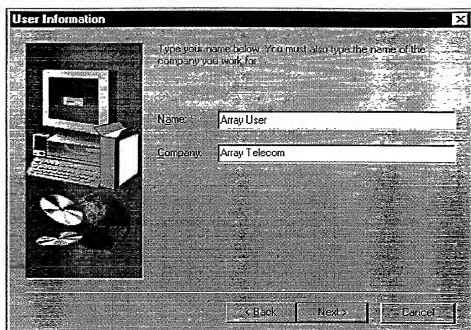


Fig. 38




Fig. 39




Fig. 40

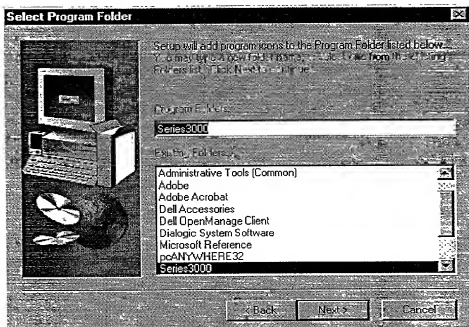


Fig. 41

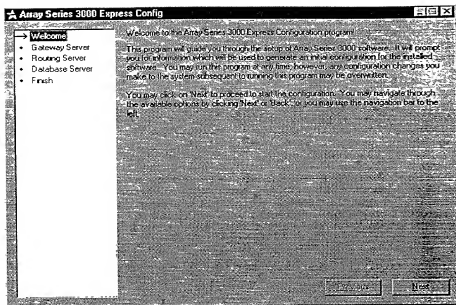


Fig. 42

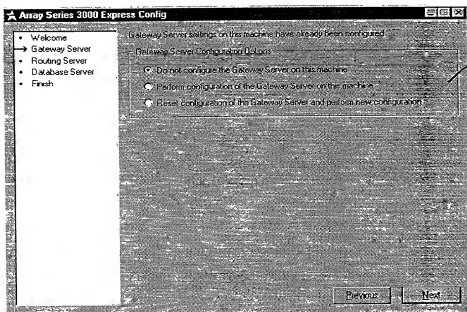


Fig. 43

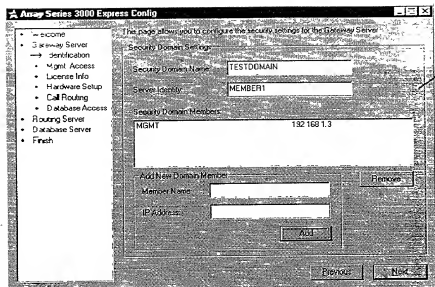


Fig. 44

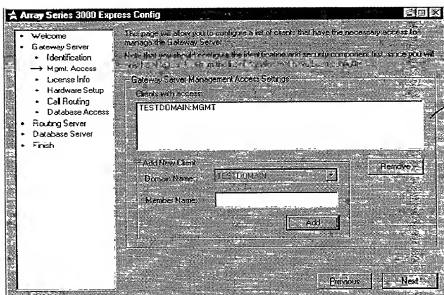


Fig. 45

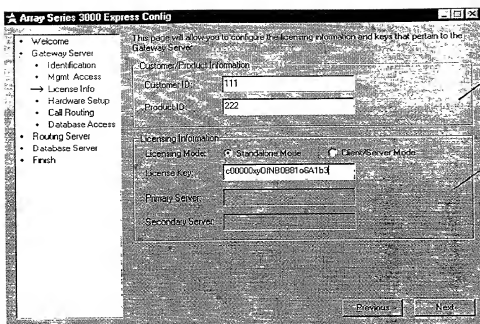


Fig. 46

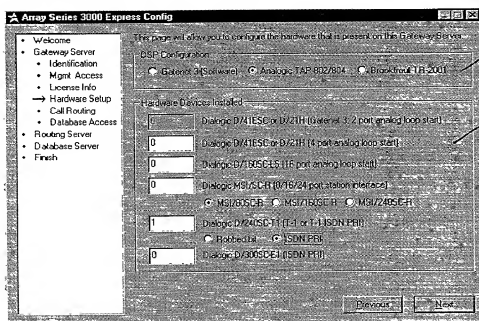


Fig. 47

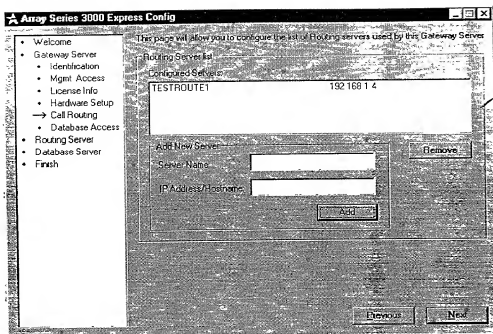


Fig. 48

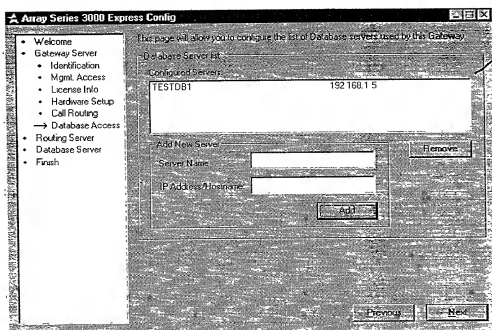


Fig. 49A

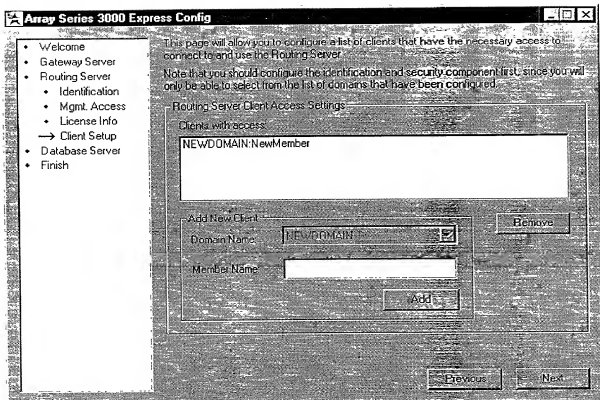


Fig. 49B

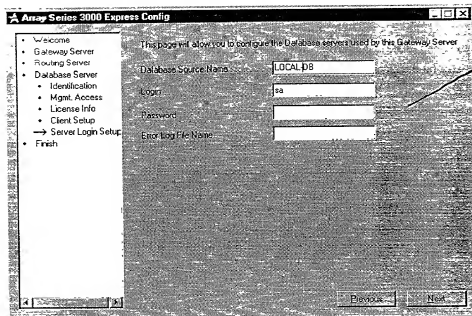


Fig. 50

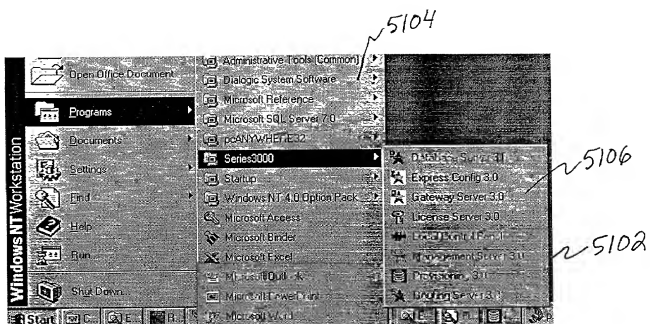


Fig. 51




Fig. 52

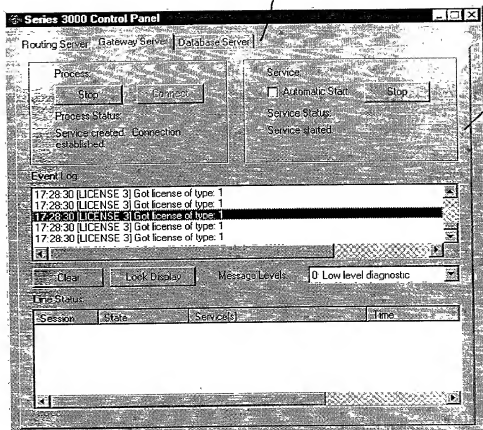


Fig. 53

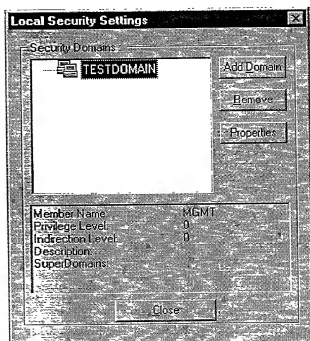


Fig. 54

Add New Managed System

S3000 Server Settings

System Name:

System Type: **Gateway Server**

IP Address:

Port Number: **6789**

Description:

OK Cancel

~5504

~5502

Fig. 55

MEMBER1

- + Line Groups
- + Services
- + Routing
- + Database Server Access
- + Hardware Devices
- + Licensing
- + Security and Identification
- + Management Access
- + General

~5602

Fig. 56

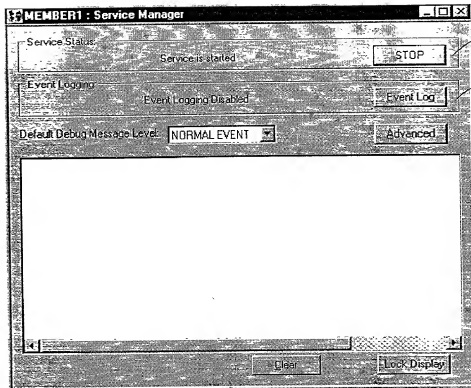


Fig. 57

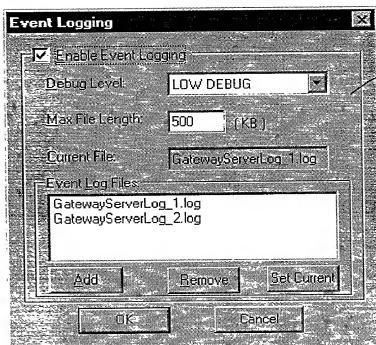


Fig. 58

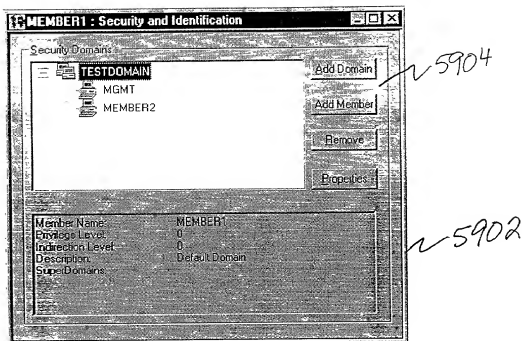


Fig. 59

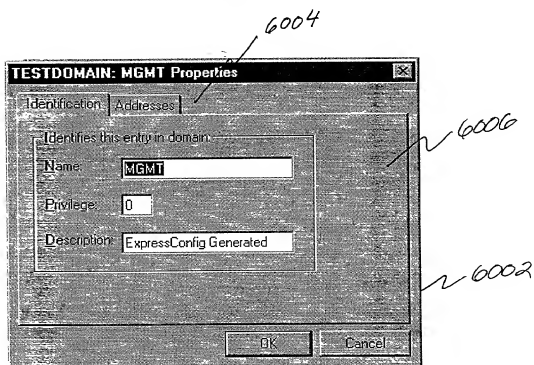


Fig. 60

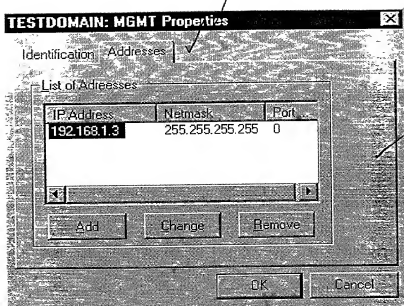


Fig. 61

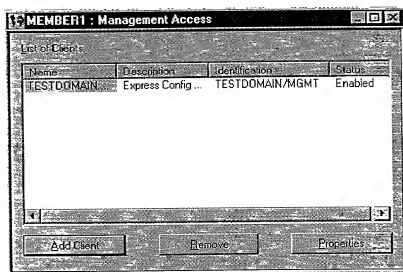


Fig. 62




Fig. 63




Fig. 64

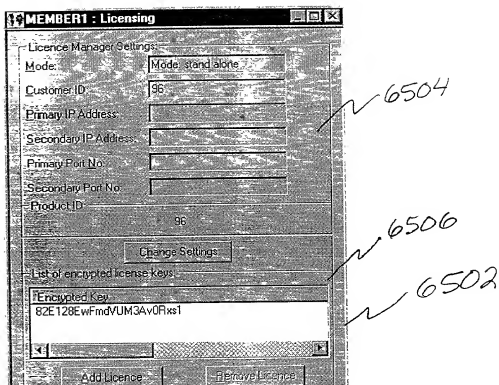


Fig. 65




Fig. 66




Fig. 67

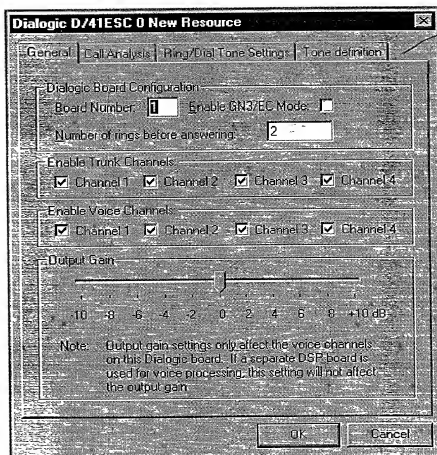


Fig. 68

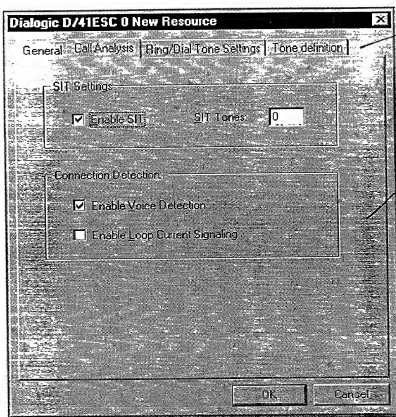


Fig. 69

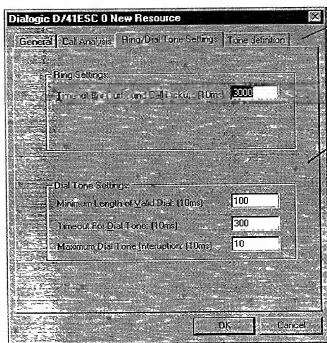


Fig. 70

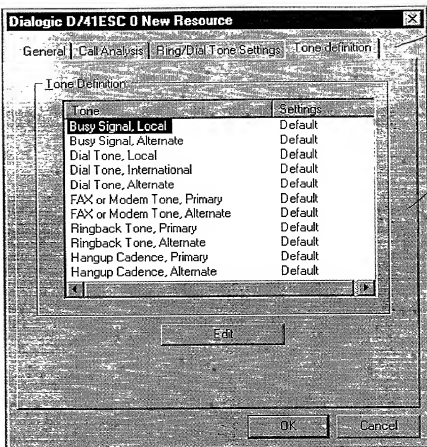


Fig. 71

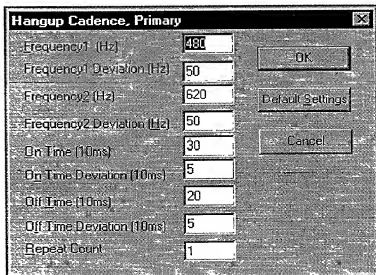


Fig. 72

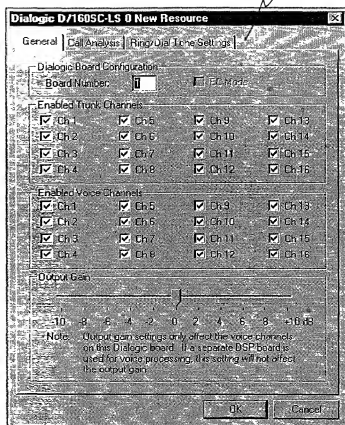


Fig. 73

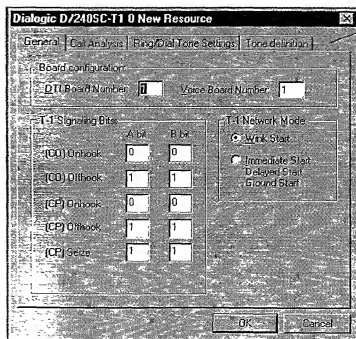


Fig. 74

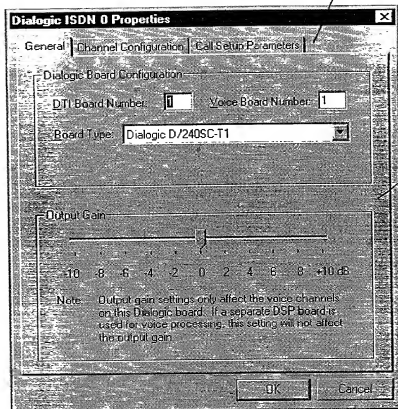


Fig. 75

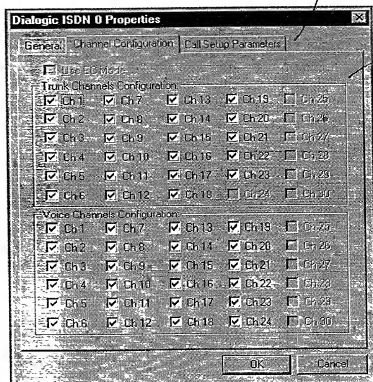


Fig. 76

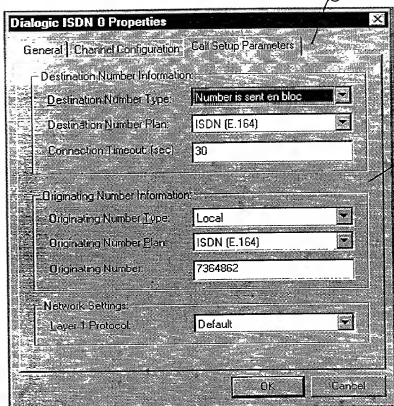


Fig. 77

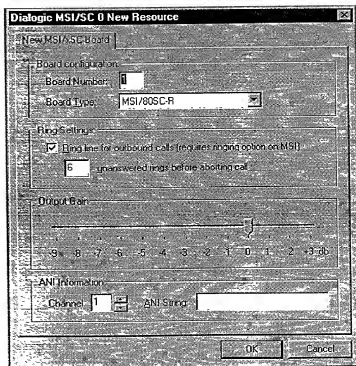


Fig. 78

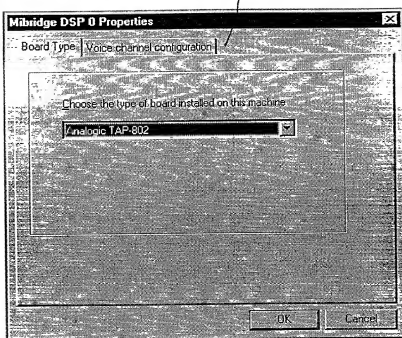


Fig. 79

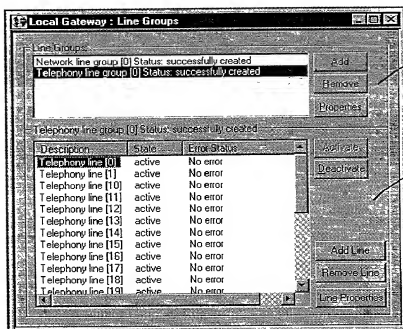


Fig. 80

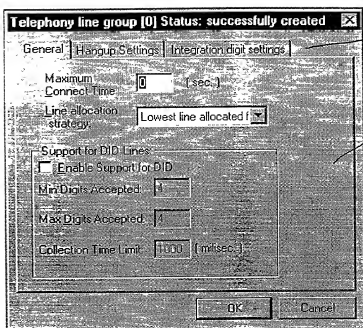


Fig. 81

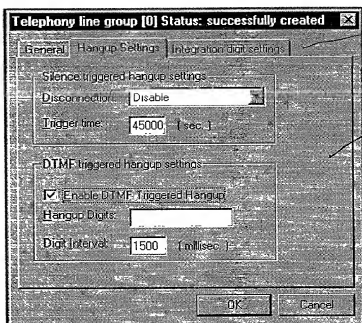


Fig. 82

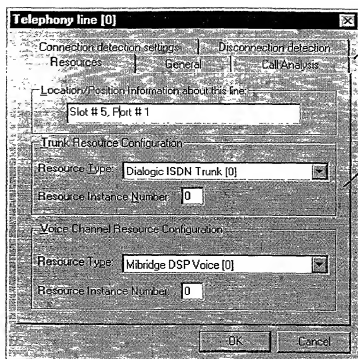


Fig. 83

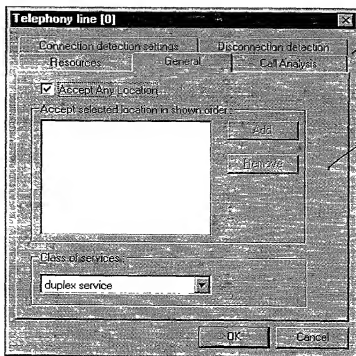


Fig. 84

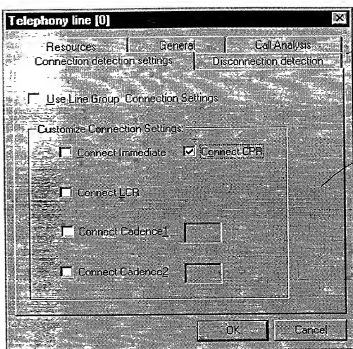


Fig. 85




Fig. 86




Fig. 87

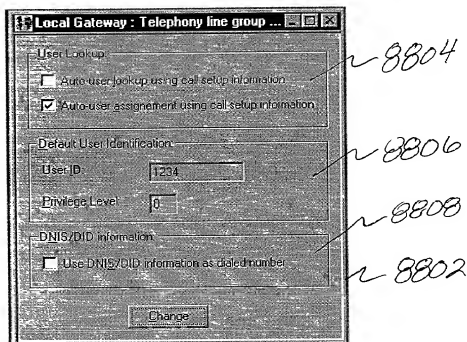


Fig. 88

Change AutoForward Service Settings

Call Address Settings

Country Code:

Area Code:

Local Number:

☒ Enable Voice Prompt

Voice Prompt:

OK Cancel

8904

8906

8902

Fig. 89




Fig. 90




Fig. 91

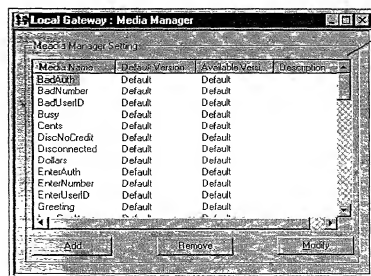


Fig. 92

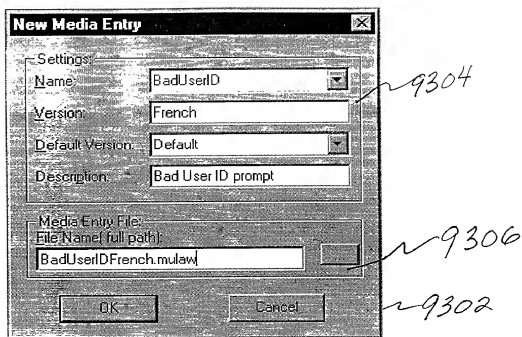


Fig. 93

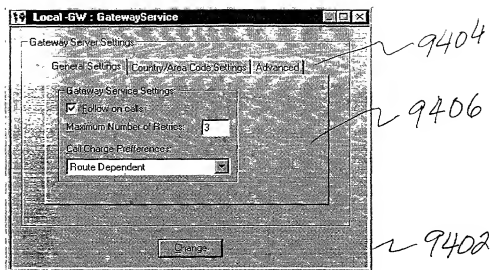


Fig. 94

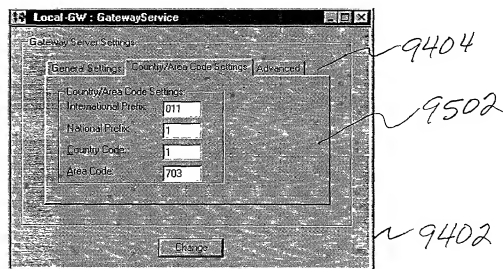


Fig. 95

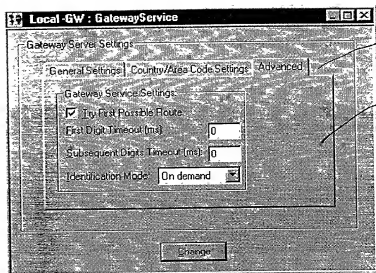


Fig. 96


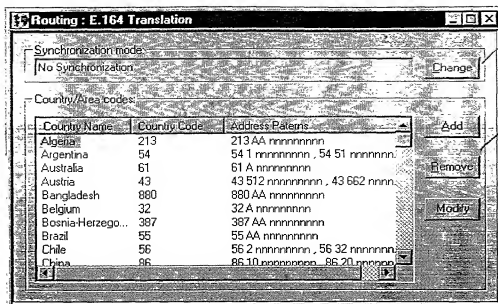


Fig. 97



10004 Fig. 99

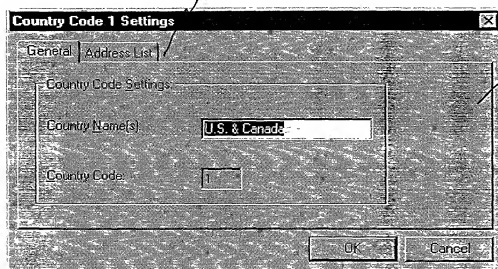


Fig. 100




Fig. 101




Fig. 102




Fig. 103




Fig. 104




Fig. 105

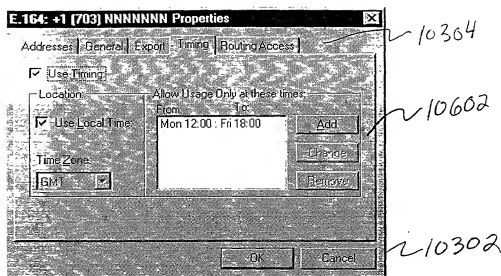


Fig. 106

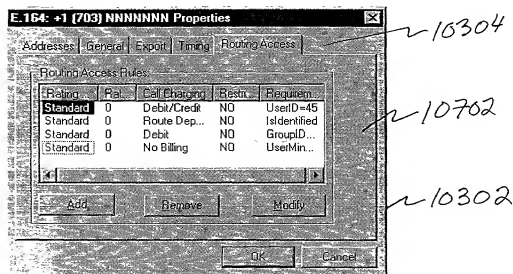


Fig. 107

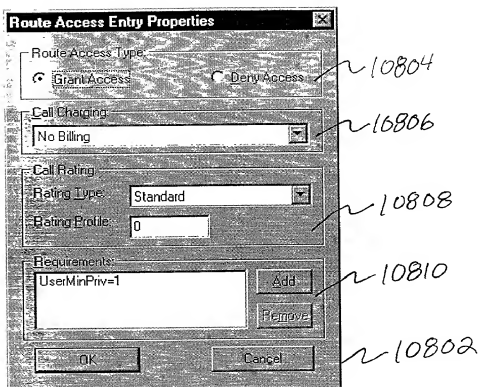


Fig. 108

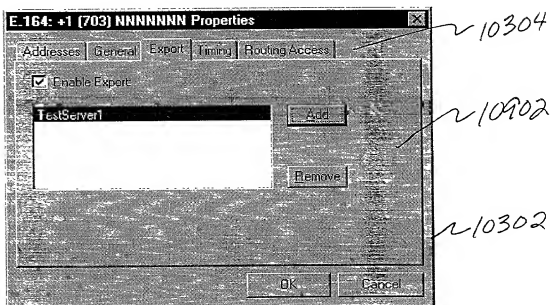


Fig. 109

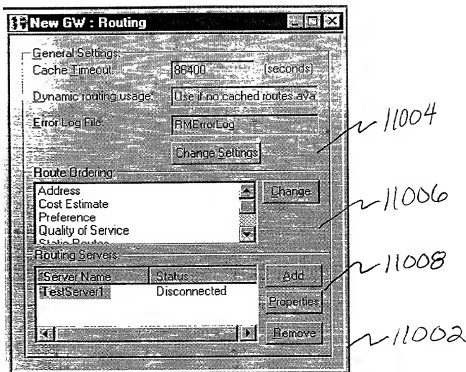


Fig. 110

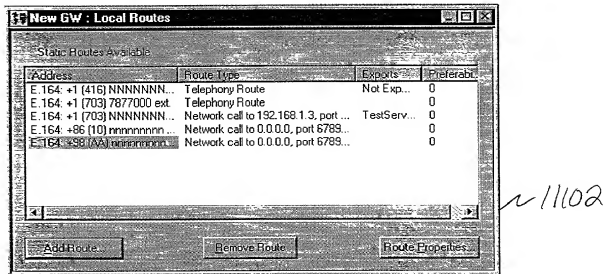


Fig. 111

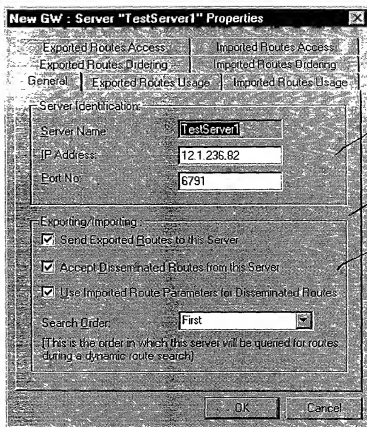


Fig. 112

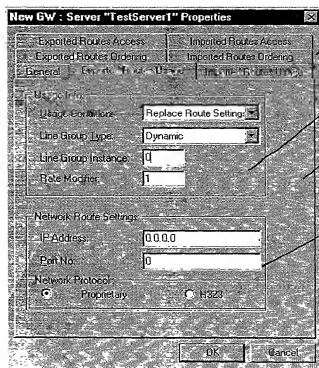


Fig. 113

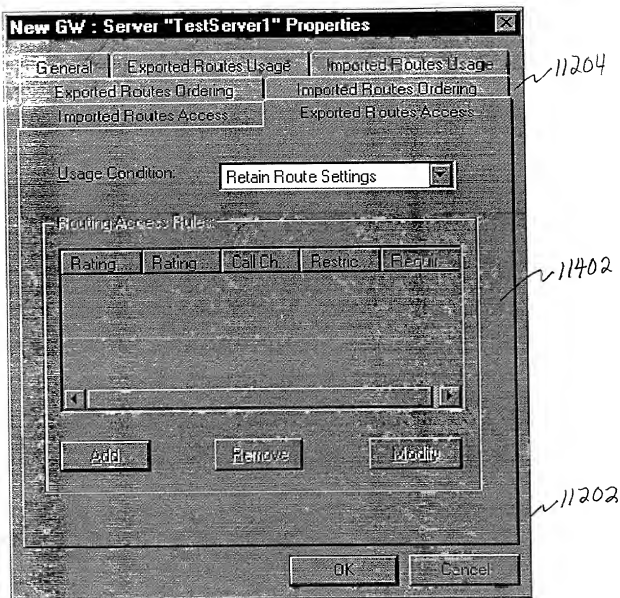


Fig. 114

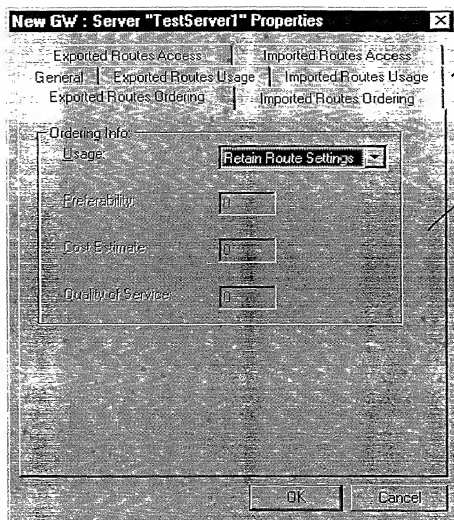


Fig. 115

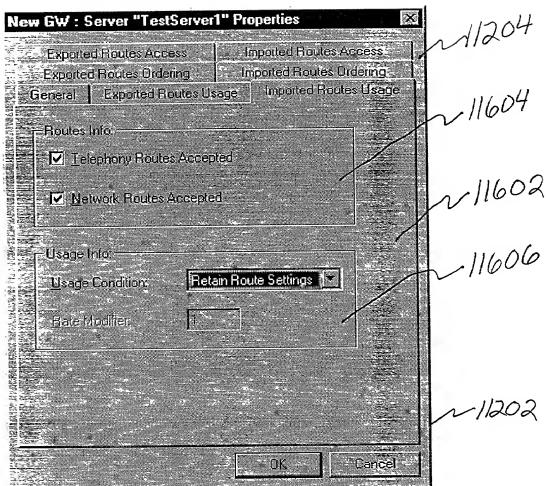


Fig. 116

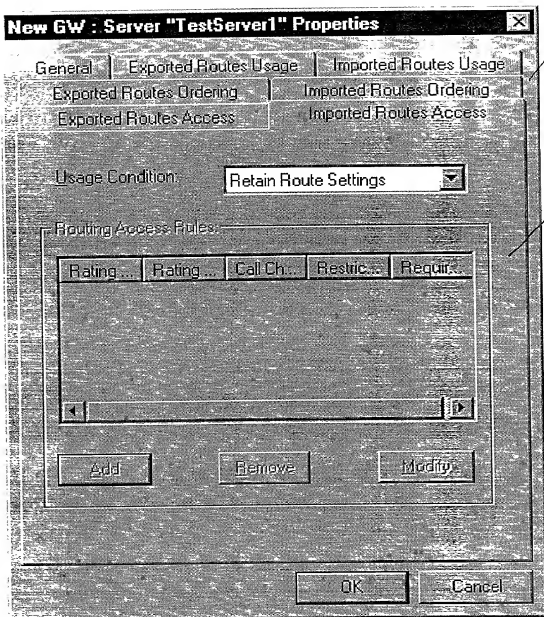


Fig. 117

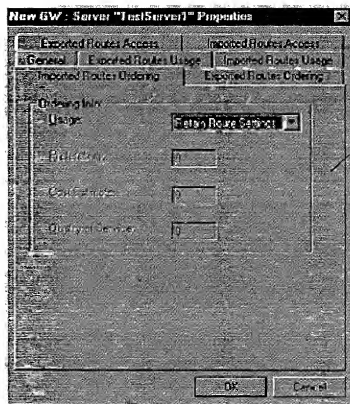


Fig. 118

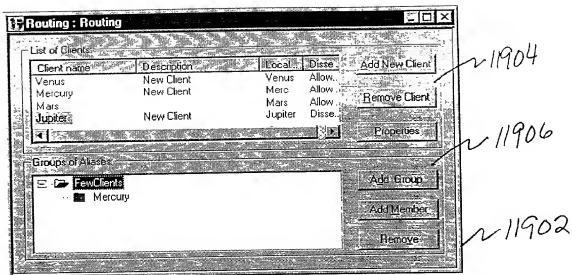


Fig. 119

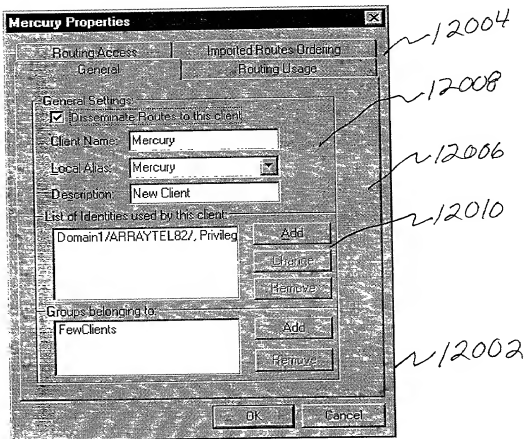


Fig. 120

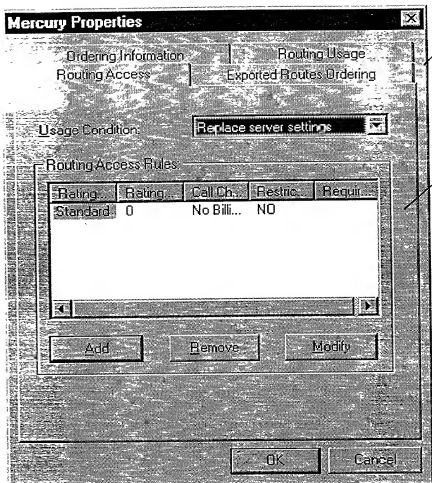


Fig. 121

Routing : Exported Routes

Exported routes

Address	Route Type	Exported by	Time Received
E.164: +1 (703) NNN...	Network call to 192.16...	Mercury	99 Aug 18 19:22
E.164: +1 (416) NNN...	Network call to 0.0.0.0...	Mercury	99 Aug 18 19:22
E.164: +1 (703) 7877...	Network call to 0.0.0.0...	Mercury	99 Aug 18 19:22

Refresh

~12202

Fig. 122

Routing : Static Routes

Static Routes Available

Address	Route Type	Prefer
E.164: +1 (333) NNNNNNN ext.	Telephony Route	0
E.164: +1 (416) NNNNNNN ext.	Telephony Route	0
E.164: +1 (416) NNNNNNN ext.	Telephony Route	0
E.164: +1 (703) 7877000 ext.	Telephony Route	0
E.164: +1 (703) nnnnnnn ext.	Telephony Route	0
E.164: +1 (777) NNNNNNN ext.	Network call to 12.1.236.82, port ...	0

Add Route... Remove Route Route Properties

~12302

Fig. 123

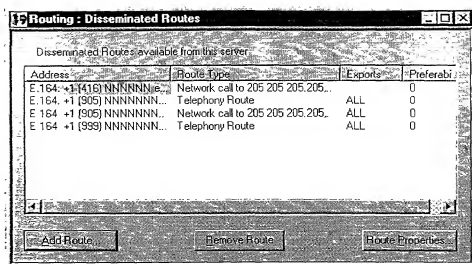


Fig. 124

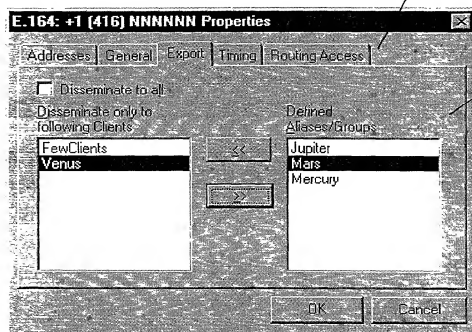


Fig. 125

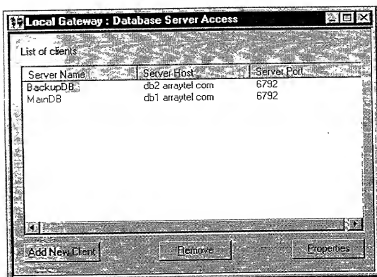


Fig. 126




Fig. 127

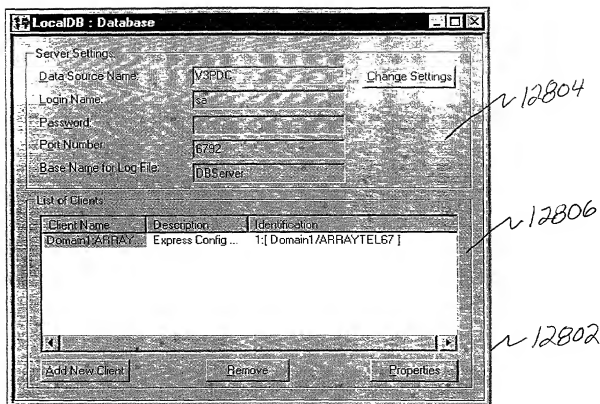


Fig. 128

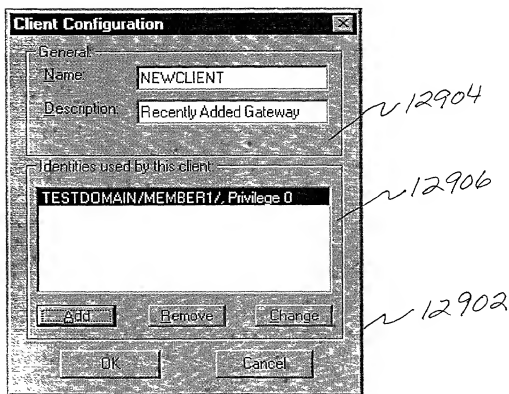


Fig. 129

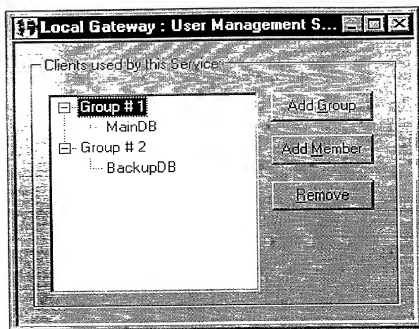


Fig. 130

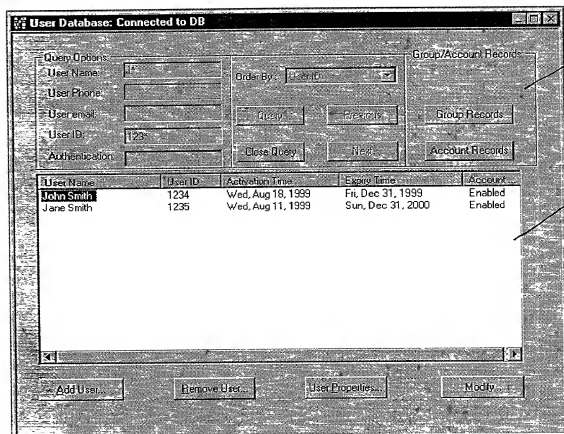


Fig. 131

Add New User

Caller ID Authentication | Account Record Status | Parent Group
 Identification | Permissions | Caller ID Ident

Identification
 Name: John Smith
 Numeric ID: 1234
 Authentication:
☒ Not Required ☐ Required ☐ Use Group
 Numeric Authentication:
 Language Pref.: English
 Email Address: John.Smith@arraytel.com
 Home phone: 1(703)555-1234
 Online Announcements:
☒ Use Group ☐ Debt Balance ☐ Time Left

OK Cancel

Fig. 132

John Smith Properties

Caller ID Authentication | Account Record Status | Parent Group
 Identification | Permissions | Caller ID Ident

Account Status:
☒ Enable ☐ Disable ☐ Use Group

Permissions:
 Normal Privilege: 1
 Authenticated Privilege: 2
 Concurrent Usage: 1

☐ Use Group Activation Time:
 YY MM DD HH:MM
 99 8 18 22:00

☐ Use Group Expiry Time:
 YY MM DD HH:MM
 99 12 31 23:00

Close

Fig. 133

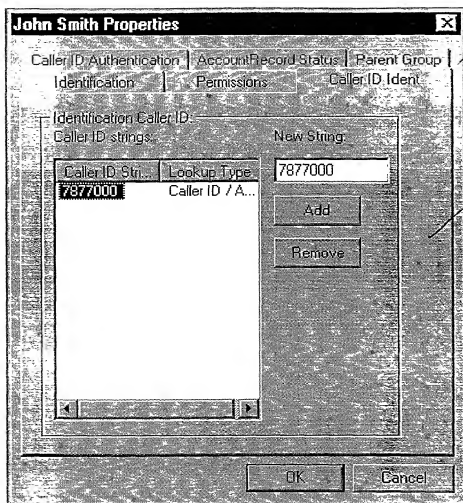


Fig. 134

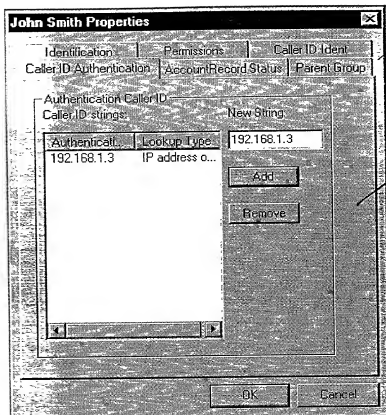


Fig. 135

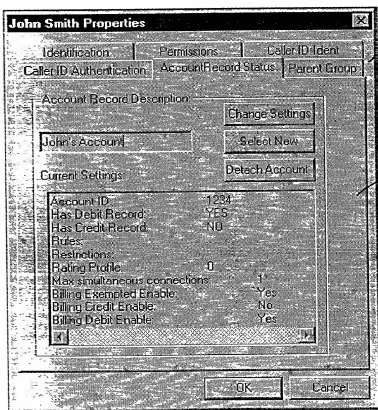


Fig. 136

John Smith Properties

Identification | Permissions | Caller ID Ident |
 Caller ID Authentication | Account Record Status | Parent Group

Group Record Description:

Current Settings

Group ID:	1234
Status:	Enabled
Start Time:	Thu, Aug 19, 1999
End Time:	Thu, Aug 19, 1999
Account ID:	1234
Max simultaneous connections:	1
Language Pref:	English
Normal Privilege:	1
Auth Privilege:	1
Description:	145566888

Fig. 137

Group Record Query

Query Options

Group ID: Order By:

Account ID:

Starting Date:

Ending Date:

Account Records

Group ID	Account ID	Starting Date	Ending Date
1234	1234	Thu, Aug 19, 1999	Thu, Aug 19, 1999

Fig. 138

13904

1234 Change Settings

General | Timing

Group ID: 1234

Account ID: 1234

Select New

Detach Account

Privilege Level: 1

☐ Force Authentication

Authentication Level: 1

Max Concurrent Usage: 1

Account Status:

☒ Enable ☐ Disable

Language Pref: English

Group Info: People named Smith

Locked By:

OK Cancel

13906

13902

Fig. 139

Account Record Query

Query Options:

Account ID:

Group ID:

User ID:

Billing Type:

Simultaneous Connections:

Order By:

Account Records:

Account ID	Debit Record	Credit Record	Billing Profile	Max Connect
1234	YES	NO	0	1
32131	YES	NO	0	1
213121	YES	NO	0	1

14004

14006

14002

Fig. 140

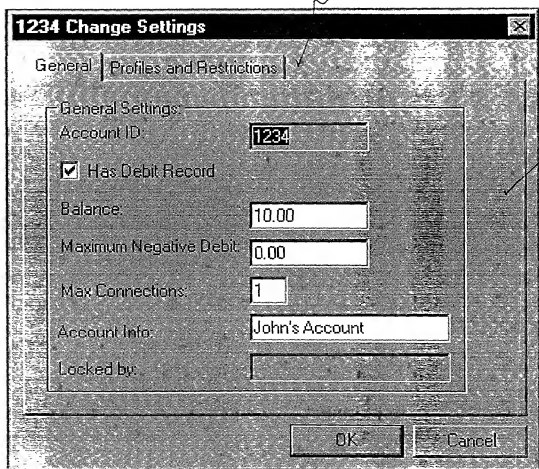


Fig. 141

260

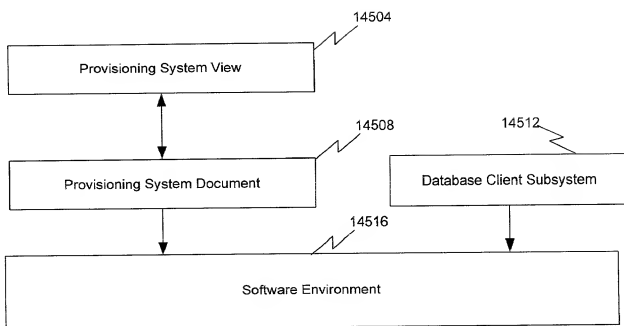



FIG. 145



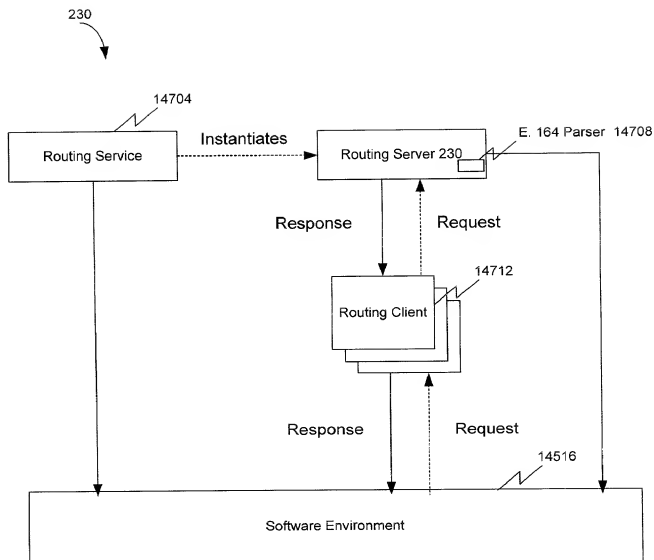


FIG. 147

452, 454, 456




FIG. 148

210

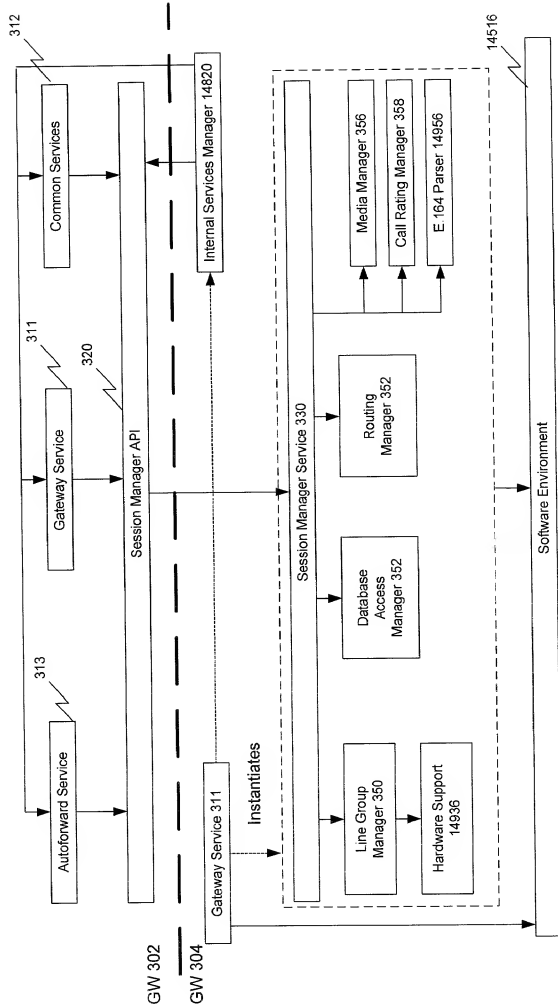


FIG. 149

15000

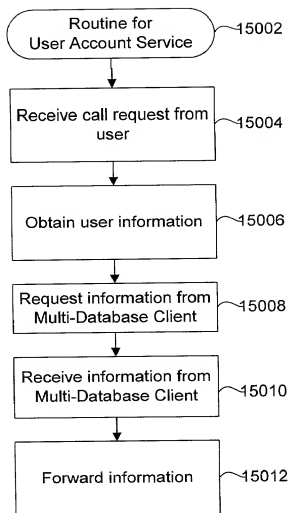


FIG. 150

15100

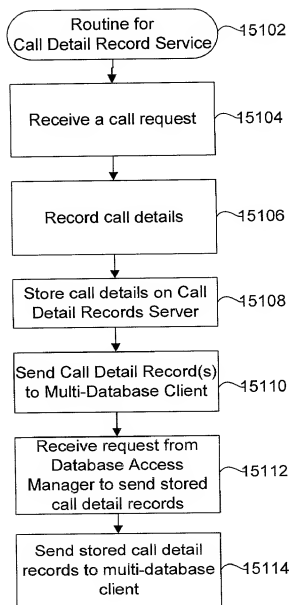


FIG. 151

15200




FIG. 152

15300

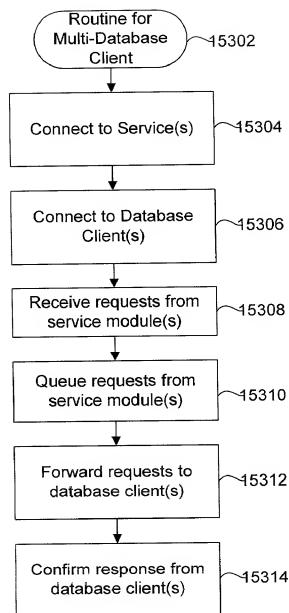


FIG. 153

15400




FIG. 154

15500




FIG. 155

15600

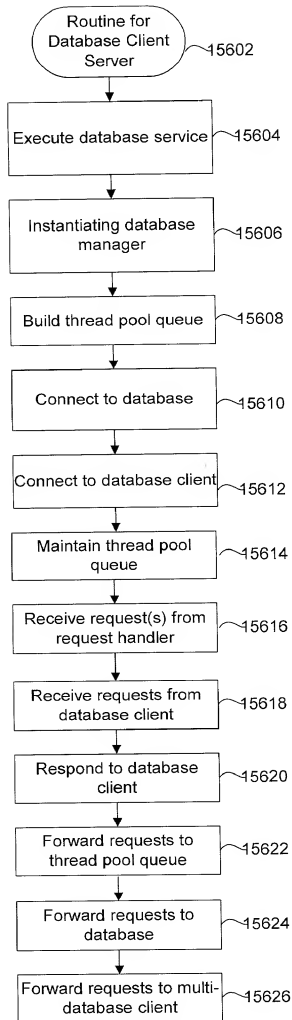


FIG. 156

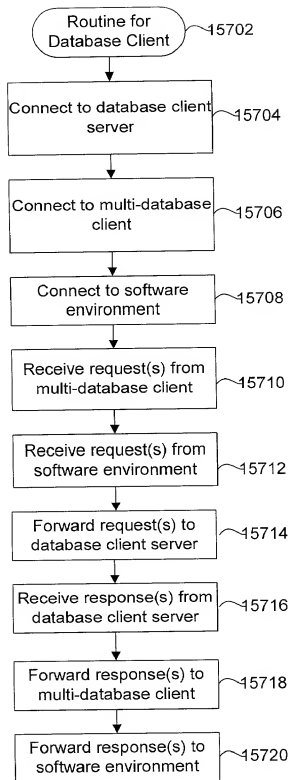


FIG. 157

15800

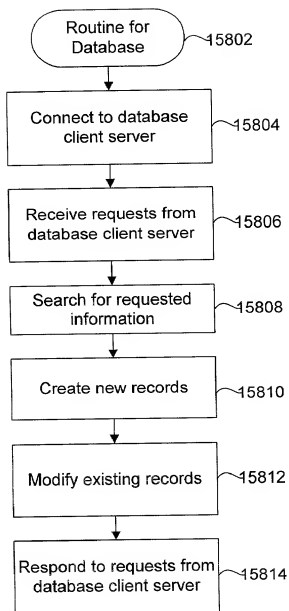


FIG. 158

15900

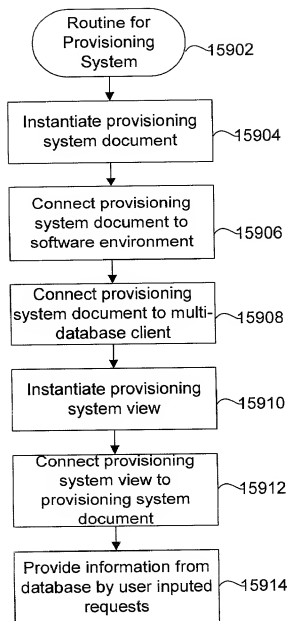


FIG. 159

16000

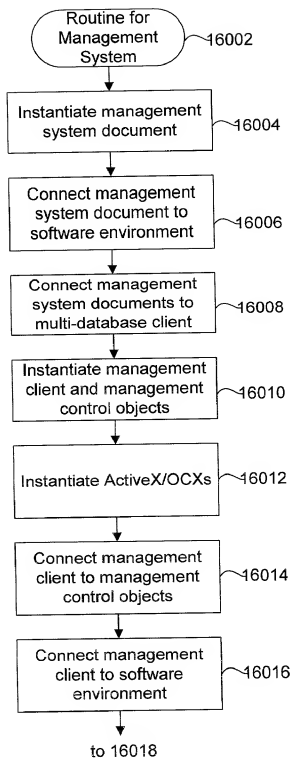


FIG. 160A

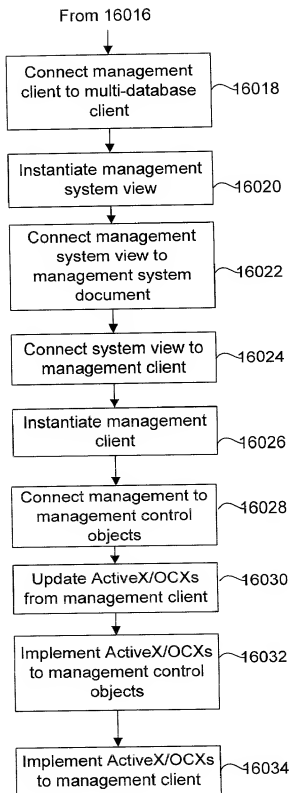


FIG. 160B

16100

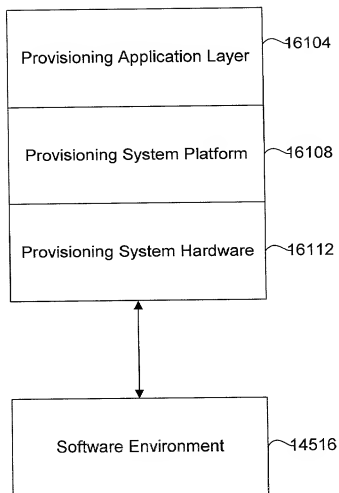


FIG. 161

16200




FIG. 162

16300

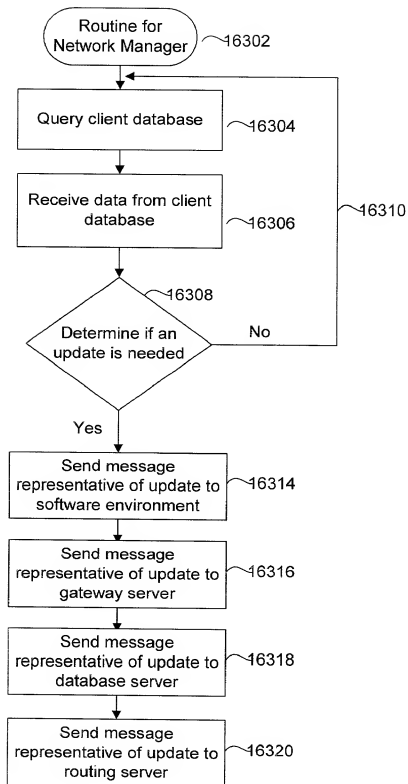


FIG. 163

16400

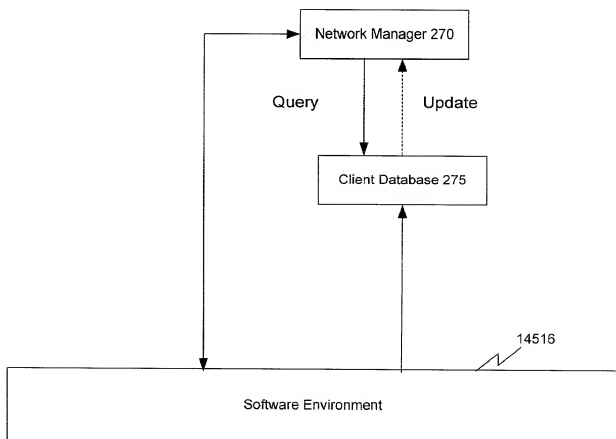


FIG. 164

Combined Declaration and Power of Attorney for Patent Application

Docket Number: **1848.0060001**

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below next to my name.

I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter that is claimed and for which a patent is sought on the invention entitled, **Method, System, and Computer Program Product for Managing Database Servers and Services**, the specification of which is attached hereto unless the following box is checked:

- ☐ was filed on _____;
as United States Application Number or PCT International Application Number _____; and
was amended on _____ (if applicable).

I hereby state that I have reviewed and understand the contents of the above identified specification, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose information that is material to patentability as defined in 37 C.F.R. § 1.56.

I hereby claim foreign priority benefits under 35 U.S.C. § 119(a)-(d) or § 365(b) of any foreign application(s) for patent or inventor's certificate, or § 365(a) of any PCT international application, which designated at least one country other than the United States listed below, and have also identified below any foreign application for patent or inventor's certificate, or PCT international application having a filing date before that of the application on which priority is claimed.

Prior Foreign Application(s)	Priority Claimed
_____ (Application No.) _____ (Country) _____ (Day/Month/Year Filed)	<input type="checkbox"/> Yes <input type="checkbox"/> No
_____ (Application No.) _____ (Country) _____ (Day/Month/Year Filed)	<input type="checkbox"/> Yes <input type="checkbox"/> No

I hereby claim the benefit under 35 U.S.C. § 119(e) of any United States provisional application(s) listed below.

_____ (Application No.)	_____ (Filing Date)
_____ (Application No.)	_____ (Filing Date)

I hereby claim the benefit under 35 U.S.C. § 120 of any United States application(s), or under § 365(c) of any PCT international application designating the United States, listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States or PCT international application in the manner provided by the first paragraph of 35 U.S.C. § 112, I acknowledge the duty to disclose information that is material to patentability as defined in 37 C.F.R. § 1.56 that became available between the filing date of the prior application and the national or PCT international filing date of this application.

09/393,658 _____ (Application No.)	September 10, 1999 _____ (Filing Date)	Pending _____ (Status - patented, pending, abandoned)
_____ (Application No.)	_____ (Filing Date)	_____ (Status - patented, pending, abandoned)

I hereby appoint the following attorney(s) and/or agent(s) to prosecute this application and to transact all business in the Patent and Trademark Office connected therewith:

Robert G. Sterne, Esq., Reg. No. 28,912; Edward J. Kessler, Esq., Reg. No. 25,688; Jorge A. Goldstein, Esq., Reg. No. 29,021; Samuel L. Fox, Esq., Reg. No. 30,353; David K.S. Cornwell, Esq., Reg. No. 31,944; Robert W. Esmond, Esq., Reg. No. 32,893; Tracy-Gene G. Durkin, Esq., Reg. No. 32,831; Michele A. Cimbala, Esq., Reg. No. 33,851; Michael B. Ray, Esq., Reg. No. 33,997; Robert E. Sokohl, Esq., Reg. No. 36,013; Eric K. Steffe, Esq., Reg. No. 36,688; Michael Q. Lee, Esq., Reg. No. 35,239; Steven R. Ludwig, Esq., Reg. No. 36,203; Raz E. Fleshner, Esq., Reg. No. 34,331; John M. Covert, Esq., Reg. No. 38,759; and Linda E. Alcorn, Esq., Reg. No. 39,588.

Send Correspondence to:

STERNE, KESSLER, GOLDSTEIN & FOX P.L.L.C.
1100 New York Avenue, N.W.
Suite 600
Washington, D.C. 20005-3934

Direct Telephone Calls to:

(202) 371-2600

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

Full name of sole or first inventor MARK SCOTT		
Signature of sole or first inventor X	X	Date
Residence 43757 Raleigh Place, Ashburn, VA 20147		
Citizenship CANADIAN		
Post Office Address Same as above		
Full name of second inventor WILLIAM WONG		
Signature of second inventor X	X	Date
Residence 67 Redstone Road, Richmond Hill, Ontario L4S 1S3, Canada		
Citizenship CANADIAN		
Post Office Address Same as above		

Full name of third inventor ANITA CHENG		
Signature of third inventor X	X	Date
Residence 43757 Raleigh Place, Ashburn, VA 20147		
Citizenship CANADIAN		
Post Office Address Same as above		
Full name of fourth inventor SIMON HO		
Signature of fourth inventor X	X	Date
Residence 43757 Raleigh Place, Ashburn, VA 20147		
Citizenship CANADIAN		
Post Office Address Same as above		
Full name of fifth inventor GEORGE IRIMESCU		
Signature of fifth inventor X	X	Date
Residence 1 Superior Avenue, Apt. 8, Etobicoke, Ontario M8V 2M1 Canada		
Citizenship CANADIAN		
Post Office Address Same as above		
Full name of sixth inventor DOREL VOINEAG		
Signature of sixth inventor X	X	Date
Residence 1505 Wilson Avenue, Apt. 302, Toronto, Ontario M3M, 1G9, Canada		
Citizenship ROMANIAN		
Post Office Address Same as above		

Full name of seventh inventor MIN YAO		
Signature of seventh inventor X	X	Date
Residence 7411 Parkwood Court, Apt. 202, Falls Church, VA 22042		
Citizenship CHINESE		
Post Office Address Same as above		
Full name of eighth inventor ROW J. ZADEH		
Signature of eighth inventor X	X	Date
Residence 1200 Olds Harrods Creek Road, Anchorage, KY 40223		
Citizenship AUSTRALIAN		
Post Office Address Same as above		

INVENTOR(S) BLADES (Last Name) (18480060001) DEC
PACF Rev. 12/2019 Page

(Supply similar information and signature for subsequent joint inventors, if any)